

Minimizing Regret in Discounted-Sum Games

Paul Hunter*, Guillermo A. Pérez†, and Jean-François Raskin*

Département d'Informatique, Université Libre de Bruxelles (ULB)
`{phunter,gperezme,jraskin}@ulb.ac.be`

April 26, 2016

Abstract

In this paper, we study the problem of minimizing regret in discounted-sum games played on weighted game graphs. We give algorithms for the general problem of computing the minimal regret of the controller (Eve) as well as several variants depending on which strategies the environment (Adam) is permitted to use. We also consider the problem of synthesizing regret-free strategies for Eve in each of these scenarios.

1 Introduction

Two-player games played by Eve and Adam on weighted graphs is a well accepted mathematical formalism for modelling quantitative aspects of a controller (Eve) interacting with its environment (Adam). The outcome of the interaction between the two players is an infinite path in the weighted graph and a value is associated to this infinite path using a *measure* such as *e.g.* the mean-payoff of the weights of edges traversed by the infinite path, or the discounted sum of those weights. In the classical model, the game is considered to be zero sum: the two players have antagonistic goals—one of the player want to maximize the value associated to the outcome while the other want to minimize this value. The main solution concept is then the notion of winning strategy and the main decision problem asks, given a threshold c , whether Eve has a strategy to ensure that, no matter how Adam plays, that the outcome has a value larger than or equal to c .

When the environment is not fully antagonistic, it is reasonable to study other solution concepts. One interesting concept to explore is the concept of *regret minimization* [3] which is as follows. When a strategy of Adam is fixed, we can identify the set of Eve's strategies that allow her to secure the *best possible outcome* against this strategy. This constitutes Eve's *best response*. Then we define the regret of a strategy σ of Eve as the difference between Eve's best response; and the payoff she secures thanks to her strategy σ . So, when trying to minimize the regret associated to a strategy, we use best responses as a *yardstick*. Let us now illustrate this with an example.

Example 1 (Investment advice). Consider the discounted sum game of Fig. 2. It models the rentability of different investment plans with a time horizon of two periods. In the first period, it can be decided to invest in treasure bonds (B) or to invest in the stock market (S). In the former case, treasure bonds (B) are chosen for two periods. In the latter case, after one period, there is again a choice for either treasure bonds (B) or stock market (S). The returns of the different investments depend on the fluctuation of the rate of interests. When the rate of interests is low (L) then the return for the stock market investments is equal to 12 and for the treasure bonds it is equal to 8. When the interest rate is high (H) then the returns for the stock market investments is equal to -4 and for the treasure bonds it is equal to -2 . To model time and take into account the inflation rate, say equal to 2 percent, we consider a discount factor $\lambda = 0.98$ for the returns. In this example, we make the hypothesis that the fluctuation of the rate of interests is *not* a function of the behavior of the investor. It means that this fluctuation rate is either one

*Authors supported by the ERC inVEST (279499) project.

†Author supported by F.R.S.-FNRS fellowship.




Figure 1: A game in which waiting is required to minimize regret.




Figure 2: A game that models different investment strategies.

	HH	HL	LH	LL	Worst-case	Regret
SS	-7.7616	7.6048	7.9784	23.2848	-7.7616	3.8808
SB	-5.8408	3.7632	9.8392	19.4432	-5.8408	3.8416
BB	-3.8808	5.7232	5.9192	15.5232	-3.8808	7.7616

Table 1: The possible rate configuration for the rate of interests are given as the first four columns, the follows the worst-case performance and the regret associated to each strategy of Eve that are given in rows. Entries in bold are the values that are maximizing the worst-case (strategy BB) and minimizing the regret (strategy SB).

of the following four possibilities: HH, HL, LH, LL. This corresponds to Adam playing a *word strategy* in our terminology. The discounted sum of returns obtained under the 12 different scenarios are given in Table 1.

Now, assume that you are a broker and you need to advise one of your customers regarding his next investment. There are several ways to advise your customer. First, if your customer is strongly *risk averse*, then you should be able to convince him that he has to go for the treasure bonds (B). Indeed, this is the choice that *maximizes the worst-case*: if the interest rates stay high for two periods (HH) then the loss will be -3.8808 while it will be higher for any other choices. Second, and maybe more interestingly, if your customer tolerates some risks, then you may want to keep him happy so that he will continue to ask for your advice in the future! Then you should propose the following strategy: first invest in the stock market (S) then in treasure bonds (B) as this strategy *minimizes regret*. Indeed, at the end of the two investment periods, the actual interest rates will be known and so your customer will evaluate your advices *ex-post*. So, after the two periods, the value of the choices made *ex ante* can be compared to the best strategy that could have been chosen knowing the evolution of the interest rates. The regret of SB is at most equal to 3.8416 in all cases and it is minimal: the regret of BB can be as high as 7.7616 if LL is observed, and the regret of SS can be as high as 3.8808.

Finally, let us remark that if the investments are done in financial markets that are subject to different interest rates, then instead of considering the minimization of regret against word strategies, then we could consider the regret against all strategies. We also study this case in this paper.

Previous works. In [9], we studied regret minimization in the context of reactive synthesis for shortest path objectives. Recently in [13], we studied the notion of regret minimization when we assume different sets of strategies from which Adam chooses. We have considered three cases: when the Adam is allowed to play *any* strategy, when he is restricted to play a *memoryless* strategy, and when he plays *word* strategies. We refer the interested reader to [13] for motivations behind each of these definitions. In that paper, we studied the regret minimization problem for the following classical quantitative measures: inf, sup, lim inf, lim sup and the *mean-payoff* measure. In this paper, we complete this picture by studying the regret minimization problem for the *discounted-sum* measure. Discounted-sum is a central measure

	Any strategy	Memoryless strategies	Word strategies
regret threshold	NP (Thm. 1)	PSPACE (Thm. 3), coNP-h (Thm. 5)	PSPACE-c (ε -gap) (Thm. 8, Thm. 9)
regret-free	PTIME (Thm. 2)	PSPACE (Thm. 4), coNP-h (Thm. 5)	NP-c (Thm. 6)

Table 2: Complexity of deciding the regret threshold and regret-free problems for fixed λ .

in quantitative games but we did not consider it in [13] because it requires specific techniques which are more involved than the ones used for the other quantitative measures. For example, while for mean-payoff objectives, strategies that minimize regret are memoryless when the Adam can play any strategy, we show in this paper that pseudo-polynomial memory is necessary (and sufficient) to minimize regret in discounted-sum games. The need for memory is illustrated by the following example.

Example 2. Consider the example in Figure 1 where $M \gg 1$. Eve can play the following strategies in this game: let $i \in \mathbb{N} \cup \{\infty\}$, and note σ^i the strategy that first plays i rounds the edge (v_I, v) and then switches to (v_I, x) . The regret values associated to those strategies are as follows.

The regret of σ^∞ is $\frac{1}{1-\lambda}$ and it is witnessed when Adam never plays the edge (v, y) . Indeed, the discounted sum of the outcome in that case is 0, while if Eve had chosen to play (v_I, x) at the first step instead, then she would have gained $\frac{1}{1-\lambda}$. The regret of σ^i is equal to the maximum between $\frac{1}{1-\lambda} - \lambda^{2i} \frac{1}{1-\lambda}$ and $\lambda^{2i+1} \frac{M}{1-\lambda} - \lambda^{2i} \frac{1}{1-\lambda}$. The maximum is either witnessed when Adam never plays (v, y) or plays (v, y) if the edge (v_I, x) has been chosen $i + 1$ times (one more time compared to σ^i).

So the strategy that minimizes regret is the strategy σ^N for $N > \frac{-\log M}{2 \log \lambda} - \frac{1}{2}$ (so that $\lambda^{2N+1} M < 1$), *i.e.* the strategy needs to count up to N .

Contributions. We describe algorithms to decide the regret threshold problem for games in three cases: when there is no restriction on the strategies that Adam can play, when Adam can only play memoryless strategies, and when Adam can only play word strategies. For this last case, our problem is closely related to open problems in the field of discounted-sum automata, and we also consider variants given as ε -gap promise problems.

We also study the complexity of the special case when the threshold is 0, *i.e.* when we ask for the existence of regret free strategies. We show that that problem is sometimes easier to solve. Our results on the complexity of both the regret threshold and the regret-free problems are summarized in Table 2. All our results are for fixed discount factor λ .

Other related works. A Boolean version of our *regret-free strategies* has been described in [7]. In that paper, they are called *remorse-free strategies*. These correspond to strategies which minimize regret in games with ω -regular objectives. They do not establish lower bounds on the complexity of realizability or synthesis of remorse-free strategies and they only consider word strategies for Adam.

In [13], we established that regret minimization when Adam plays word strategies only is a generalization of the notion of *good-for-games automata* [11] and *determinization by pruning* (of a refinement) [1].

The notion of regret is closely related to the notion of competitive ratios used for the analysis of online algorithms [14]: the performance of an online algorithm facing uncertainty (*e.g.* about the future incoming requests or data) is compared to the performance of an offline algorithm (where uncertainty is resolved). According to this quality measure, an online algorithm is better if its performance is closer to the performance of an optimal offline solution.

Structure of the paper. In Sect. 2, we introduce the necessary definitions and notations. In Sect. 3, we study the minimization of regret when the second player plays any strategy. Finally, in Sect. 4, we study the minimization of regret when the second player plays a memoryless strategy and in Sect. 5 when he plays a word strategy.

2 Preliminaries

A *weighted arena* is a tuple $G = (V, V_\exists, E, w, v_I)$ where (V, E, w) is an edge-weighted graph (with rational weights), $V_\exists \subseteq V$, and $v_I \in V$ is the initial vertex. For a given $v \in V$ we denote by $\text{succ}(v)$ the set of *successors of v in G* , that is the set $\{v \in V : (u, v) \in E\}$. We assume w.l.o.g. that no vertex is a sink, i.e. $\forall v \in V : |\text{succ}(v)| > 0$, and that every Eve vertex has more than one successor, i.e. $\forall v \in V_\exists : |\text{succ}(v)| > 1$. In the sequel, we depict vertices in V_\exists with squares and vertices in $V \setminus V_\exists$ with circles. We denote the maximum absolute value of a weight in a weighted arena by W .

A *play* in a weighted arena is an infinite sequence of vertices $\pi = v_0 v_1 \dots$ where $(v_i, v_{i+1}) \in E$ for all i . Given a play $\pi = v_0 v_1 \dots$ and integers k, l we define $\pi[k..l] := v_k \dots v_l$, $\pi[..k] := \pi[0..k]$, and $\pi[l..] := v_l v_{l+1} \dots$, all of which we refer to as *play prefixes*. To improve readability, we try to adhere to the following convention: use π to denote plays and ρ for play prefixes. The *length of a play* π , denoted $|\pi|$, is ∞ , and the *length of a play prefix* $\rho = v_0 \dots v_n$, i.e. $|\rho|$, is $n + 1$.

A *strategy for Eve (Adam)* is a function σ that maps play prefixes ending with a vertex v from V_\exists ($V \setminus V_\exists$) to a successor of v . A strategy has memory m if it can be realized as the output of a finite state machine with m states (see e.g. [12] for a formal definition). A *memoryless (or positional) strategy* is a strategy with memory 1, that is, a function that only depends on the last element of the given partial play. A play $\pi = v_0 v_1 \dots$ is *consistent with a strategy σ for Eve (Adam)* if whenever $v_i \in V_\exists$ ($v_i \in V \setminus V_\exists$), then $\sigma(\pi[..i]) = v_{i+1}$. We denote by $\Sigma_\exists(G)$ ($\Sigma_\forall(G)$) the set of all strategies for Eve (Adam) and by $\Sigma_\exists^m(G)$ ($\Sigma_\forall^m(G)$) the set of all strategies for Eve (Adam) in G that require memory of size at most m , in particular $\Sigma_\exists^1(G)$ ($\Sigma_\forall^1(G)$) is the set of all memoryless strategies for Eve (Adam) in G . We omit G if the context is clear.

Given strategies σ, τ , for Eve and Adam respectively, and $v \in V$, we denote by $\pi_{\sigma\tau}^v$ the unique play starting from v that is consistent with σ and τ . If v is omitted, it is assumed to be v_I .

A *weighted automaton* is a tuple $\Gamma = (Q, q_I, A, \Delta, w)$ where A is a finite alphabet, Q is a finite set of states, q_I is the initial state, $\Delta \subseteq Q \times A \times Q$ is the transition relation, $w : \Delta \rightarrow \mathbb{Q}$ assigns weights to transitions. A *run* of Γ on a word $a_0 a_1 \dots \in A^\omega$ is a sequence $\rho = q_0 a_0 q_1 a_1 \dots \in (Q \times A)^\omega$ such that $(q_i, a_i, q_{i+1}) \in \Delta$, for all $i \geq 0$, and has *value* $\text{Val}(\rho)$ determined by the sequence of weights of the transitions of the run and the payoff function Val . The value Γ assigns to a word w , $\Gamma(w)$, is the supremum of the values of all runs on the word. We say the automaton is deterministic if Δ is functional.

Safety games. A *safety game* is played on a non-weighted arena by Eve and Adam. The goal of Eve is to perpetually avoid traversing edges from a set of *bad edges*, while Adam attempts to force the play through any unsafe edge. More formally, a safety game is a tuple (G, B) where $G = (V, V_\exists, E, v_I)$ is a non-weighted arena and $B \subseteq E$ is the set of bad edges. A play $\pi = v_0 v_1 \dots$ is *winning for Eve* if $(v_i, v_{i+1}) \notin B$, for all $i \geq 0$, and it is *winning for Adam* otherwise. A strategy for Eve (Adam) is winning for her (him) in the safety game if all plays consistent with it are winning for her (him). A player wins the safety game if (s)he has a winning strategy.

Lemma 1 (from [2]). *Safety games are positionally determined: either Eve has a positional winning strategy or Adam has a positional strategy. Determining the winner in a safety game is decidable in linear time.*

Discounted-sum. A play in a weighted arena, or a run in a weighted automaton, induces an infinite sequence of weights. We define below the *discounted-sum* payoff function which maps finite and infinite sequences of rational weights to real numbers. In the sequel we refer to a weighted arena together with a payoff function as a *game*. Formally, given a sequence of weights $\chi = x_0 x_1 \dots$ of length $n \in \mathbb{N} \cup \{\infty\}$, the *discounted-sum* is defined by a rational discount factor $\lambda \in (0, 1)$: $\text{DS}_\lambda(\chi) := \sum_{i=0}^n \lambda^i x_i$. For convenience, we apply payoff functions directly to plays, runs, and prefixes. For instance, given a play or play prefix $\pi = v_0 v_1 \dots$ we write $\text{DS}_\lambda(\pi)$ instead $\text{DS}_\lambda(w(v_0, v_1)w(v_1, v_2) \dots)$.

Consider a fixed weighted arena G , and a discounted-sum payoff function $\text{Val} = \text{DS}_\lambda$ for some $\lambda \in (0, 1)$. Given strategies σ, τ , for Eve and Adam respectively, and $v \in V$, we denote the value of $\pi_{\sigma\tau}^v$ by $\text{Val}_G^v(\sigma, \tau) := \text{Val}(\pi_{\sigma\tau}^v)$. We omit G if it is clear from the context. If v is omitted, it is assumed to be v_I .

Antagonistic & co-operative values. Two values associated with a weighted arena that we will use throughout are the *antagonistic and co-operative values*, defined for plays from a vertex $v \in V$ as:

$$\mathbf{aVal}^v(G) := \sup_{\sigma \in \mathfrak{S}_{\exists}} \inf_{\tau \in \mathfrak{S}_{\forall}} \mathbf{Val}^v(\sigma, \tau) \quad \mathbf{cVal}^v(G) := \sup_{\sigma \in \mathfrak{S}_{\exists}} \sup_{\tau \in \mathfrak{S}_{\forall}} \mathbf{Val}^v(\sigma, \tau).$$

Again, if G is clear from the context it will be omitted, and if v is omitted it is assumed to be v_I . We note that, as memoryless strategies are sufficient in discounted-sum games [15], \mathbf{aVal} can be computed in time polynomial (in $\frac{1}{1-\lambda}$, $|V|$, and $\log_2 W$). If λ is given as part of the input, this becomes exponential (in the size of the input). Regardless of whether λ is part of the input, \mathbf{cVal} is computable in polynomial time, determining if \mathbf{aVal} is bigger (or smaller) than a given threshold is decidable and in $\mathbf{NP} \cap \mathbf{coNP}$, and the values \mathbf{cVal} and \mathbf{aVal} are representable using a polynomial number of bits.

A useful observation used by Zwick and Paterson in [15], and which is implicitly used throughout this work, is the following.

Remark 1. For all $u \in V$, $\mathbf{cVal}^u(G) = \max\{w(u, v) + \lambda \mathbf{cVal}^v(G) : (u, v) \in E\}$. For all $u \in V_{\exists}$, $\mathbf{aVal}^u(G) = \max\{w(u, v) + \lambda \mathbf{aVal}^v(G) : (u, v) \in E\}$. For all $u \in V \setminus V_{\exists}$, $\mathbf{aVal}^u(G) = \min\{w(u, v) + \lambda \mathbf{aVal}^v(G) : (u, v) \in E\}$.

We say a strategy σ for Eve is *worst-case optimal (maximizing)* from $v \in V$ if it holds that $\inf_{\tau \in \mathfrak{S}_{\forall}} \mathbf{Val}^v(\sigma, \tau) = \mathbf{aVal}^v(G)$. Similarly, a strategy τ for Adam is *worst-case optimal (minimizing)* from $v \in V$ if it holds that $\sup_{\sigma \in \mathfrak{S}_{\exists}} \mathbf{Val}^v(\sigma, \tau) = \mathbf{aVal}^v(G)$. Also, a pair of strategies σ, τ for Eve and Adam, respectively, is said to be *co-operative optimal* from $v \in V$ if $\mathbf{Val}^v(\sigma, \tau) = \mathbf{cVal}^v(G)$.

Lemma 2 (from [15]). *The following hold:*

- there exists $\sigma \in \mathfrak{S}_{\exists}$ which is worst-case optimal maximizing from all $v \in V$,
- there exists $\tau \in \mathfrak{S}_{\forall}$ which is worst-case optimal minimizing from all $v \in V$,
- there are $\sigma \in \mathfrak{S}_{\exists}$ and $\tau \in \mathfrak{S}_{\forall}$ which are co-operative optimal from all $v \in V$.

We now recall the definition of a *strongly co-operative optimal* strategy σ for Eve. Formally, for any play prefix $\rho = v_0 \dots v_n$ consistent with σ , and such that $v_n \in V_{\exists}$ if $\sigma(\rho) = v'$, then $v' \in \mathbf{cOpt}(v_n)$; where $\mathbf{cOpt}(u) := \{v \in V : (u, v) \in E \text{ and } \mathbf{cVal}^u(G) = w(u, v) + \lambda \mathbf{cVal}^v(G)\}$. Finally, we define a new type of strategy for Eve: *co-operative worst-case optimal* strategies. A strategy is of this type if, for any play prefix $\rho = v_0 \dots v_n$ consistent with σ , and such that $v_n \in V_{\exists}$, if $\sigma(\rho) = v'$ then $v' \in \mathbf{wOpt}(v_n)$ and

$$w(v_n, v') + \lambda \mathbf{cVal}^{v'}(G) = \max\{w(v_n, v'') + \lambda \mathbf{cVal}^{v''}(G) : v'' \in \mathbf{wOpt}(v_n)\},$$

where $\mathbf{wOpt}(u) := \{v \in V : (u, v) \in E \text{ and } \mathbf{aVal}^u(G) = w(u, v) + \lambda \mathbf{aVal}^v(G)\}$.

It is not hard to verify that strategies of the above types always exist for Eve.

Lemma 3. *There exist strongly co-operative optimal strategies and co-operative worst-case optimal strategies for Eve.*

Regret. Let $\Sigma_{\exists} \subseteq \mathfrak{S}_{\exists}$ and $\Sigma_{\forall} \subseteq \mathfrak{S}_{\forall}$ be sets of strategies for Eve and Adam respectively. Given $\sigma \in \Sigma_{\exists}$ we define the *regret of σ in G w.r.t. Σ_{\exists} and Σ_{\forall}* as:

$$\mathbf{reg}_{\Sigma_{\exists}, \Sigma_{\forall}}^{\sigma}(G) := \sup_{\tau \in \Sigma_{\forall}} (\sup_{\sigma' \in \Sigma_{\exists}} \mathbf{Val}(\sigma', \tau) - \mathbf{Val}(\sigma, \tau)).$$

A strategy σ for Eve is then said to be *regret-free* w.r.t. Σ_{\exists} and Σ_{\forall} if $\mathbf{reg}_{\Sigma_{\exists}, \Sigma_{\forall}}^{\sigma}(G) = 0$. We define the *regret of G w.r.t. Σ_{\exists} and Σ_{\forall}* as:

$$\mathbf{Reg}_{\Sigma_{\exists}, \Sigma_{\forall}}(G) := \inf_{\sigma \in \Sigma_{\exists}} \mathbf{reg}_{\Sigma_{\exists}, \Sigma_{\forall}}^{\sigma}(G).$$

When Σ_{\exists} or Σ_{\forall} are omitted from $\mathbf{reg}(\cdot)$ and $\mathbf{Reg}(\cdot)$ they are assumed to be the set of all strategies for Eve and Adam.

In the unfolded definition of the regret of a game, *i.e.*

$$\mathbf{Reg}_{\Sigma_{\exists}, \Sigma_{\forall}}(G) := \inf_{\sigma \in \Sigma_{\exists}} \sup_{\tau \in \Sigma_{\forall}} (\sup_{\sigma' \in \Sigma_{\exists}} \mathbf{Val}(\sigma', \tau) - \mathbf{Val}(\sigma, \tau)),$$

let us refer to the witnesses σ and σ' as the *primary strategy* and the *alternative strategy* respectively. Observe that for any primary strategy for Eve and any one strategy for Adam, we can assume Adam plays to maximize the payoff (*i.e.* co-operates) against the alternative strategy once it deviates (necessarily at an Eve vertex) or to minimize against the primary strategy—again, once it deviates. Indeed, since the deviation yields different histories, the two strategies for Adam can be combined without conflict. More formally,

Lemma 4. *Consider any $\sigma \in \mathfrak{S}_\exists$, $\tau \in \mathfrak{S}_\forall$, and corresponding play $\pi_{\sigma\tau} = v_0v_1 \dots$. For all $i \geq 0$ such that $v_i \in V_\exists$, for all $v' \in \text{succ}(v_i) \setminus \{v_{i+1}\}$ there exist $\sigma' \in \mathfrak{S}_\exists$, $\tau' \in \mathfrak{S}_\forall$ for which (i) $\pi_{\sigma'\tau}[..i+1] = \pi_{\sigma\tau}[..i] \cdot v'$, (ii) $\text{Val}(\pi_{\sigma'\tau}[i+1..]) = \mathbf{cVal}^{v'}(G)$, and (iii) $\pi_{\sigma\tau} = \pi_{\sigma'\tau}$.*

Lemma 5. *Consider any $\sigma \in \mathfrak{S}_\exists$, $\tau \in \mathfrak{S}_\forall$, and corresponding play $\pi_{\sigma\tau} = v_0v_1 \dots$. For all $i \geq 0$ such that $v_i \in V_\exists$, for all $v' \in \text{succ}(v_i) \setminus \{v_{i+1}\}$ there exist $\sigma' \in \mathfrak{S}_\exists$, $\tau' \in \mathfrak{S}_\forall$ for which (i) $\pi_{\sigma'\tau}[..i+1] = \pi_{\sigma\tau}[..i] \cdot v' = \pi_{\sigma'\tau}[..i] \cdot v'$, (ii) $\text{Val}(\pi_{\sigma'\tau}[i+1..]) = \mathbf{cVal}^{v'}(G)$, and (iii) $\text{Val}(\pi_{\sigma\tau}[i+1..]) \leq \mathbf{aVal}^{v_{i+1}}(G)$.*

Both claims follow from the definitions of strategies for Eve and Adam and from Lemma 2.

In the remaining of this work, we will assume that λ is **not** given as part of the input.

3 Regret against all strategies of Adam

In this section we describe an algorithm to compute the (minimal) regret of a discounted-sum game when there are no restrictions placed on the strategies of Adam. The algorithm can be implemented by an alternating machine guaranteed to halt in polynomial time. We show that the regret *value* of any game is achieved by a strategy for Eve which consists of two strategies, the first choosing edges which lead to the optimal co-operative value, the second choosing edges which ensure the antagonistic value. The switch from the former to the latter is done based on the “local regret” of the vertex (this is formalized in the sequel). The latter allows us to claim NP-membership of the regret threshold problem. The following theorem summarizes the bounds we obtain:

Theorem 1. *Deciding if the regret value is less than a given threshold (strictly or non-strictly), playing against all strategies of Adam, is in NP.*

Let us start by formalizing the concept of *local regret*. Given a play or play prefix $\pi = v_0 \dots$ and integer $0 \leq i < |\pi|$ such that $v_i \in V_\exists$, define $\mathbf{locreg}(\pi, i)$ as follows:

$$\begin{cases} \lambda^i \left(\mathbf{cVal}_{v_{i+1}}^{v_i}(G) - \text{Val}(\pi[i..]) \right) & \text{if } \pi \text{ is a play,} \\ \lambda^i \left(\mathbf{cVal}_{v_{i+1}}^{v_i}(G) - \text{Val}(\pi[i..j]) \right) - \lambda^j \mathbf{aVal}^{v_j}(G) & \text{if } \pi \text{ is a prefix of length } j+1 > i+1, \\ \lambda^i (\mathbf{cVal}^{v_i}(G) - \mathbf{aVal}^{v_i}(G)) & \text{if } \pi \text{ is a prefix of length } i+1, \end{cases}$$

where $\mathbf{cVal}_{v_{i+1}}^{v_i}(G) = \max\{w(v_i, v) + \lambda \mathbf{cVal}^v(G) : (v_i, v) \in E \text{ and } v \neq v_{i+1}\}$. Intuitively, for π a play, $\mathbf{locreg}(\pi, i)$ corresponds to the difference between the value of the best *deviation* from position i and the value of π . For π a play prefix, $\mathbf{locreg}(\pi, i)$ assumes that after position $j = |\pi| - 1$ Eve will play a worst-case optimal strategy.

Deciding 0-regret. We will now argue that the problem of determining whether Eve has a regret-free strategy can be decided in polynomial time. Furthermore, if no such strategy for Eve exists, we will extract a strategy for Adam which, against any strategy of Eve, ensures non-zero regret. To do so, we will reduce the problem to that of deciding whether Eve wins a safety game. The unsafe edges are determined by a function of the antagonistic and co-operative values of the original game. Critically, the game is played on the same arena as the original regret game.

Theorem 2. *Deciding if the regret value is 0, playing against all strategies of Adam, is in PTIME.*

Proof. We define a partition of the edges leaving vertices from V_\exists into good and bad for Eve. A bad edge is one which witnesses non-zero local regret. We then show that Eve can ensure a regret value of 0 if and only if she has a strategy to avoid ever traversing bad edges. More formally, let us assume a given




Figure 3: Depiction of a play and a “better alternative play”.




Figure 4: A deviation from v_k cannot be a best alternative to $\pi_{\sigma\tau}$ if $j \geq N(\mathbf{Val}(\sigma', \tau) - \mathbf{Val}(\sigma, \tau))$.

weighted arena $G = (V, V_\exists, v_I, E, w)$ and a discount factor $\lambda \in (0, 1)$. We define the set of bad edges $\mathcal{B} := \{(u, v) \in E : u \in V_\exists \text{ and } w(u, v) + \lambda \mathbf{aVal}^v(G) < \mathbf{cVal}_{\neg v}^u(G)\}$.

Note that strategies for either player in the newly defined safety game are also strategies for them in the original game (and vice versa as well). We now claim that winning strategies for Adam in the safety game $\hat{G} = (V, V_\exists, v_I, E, \mathcal{B})$ ensure that, regardless of the strategy of Eve, its regret will be strictly positive. The idea behind the claim is that, Adam can force to traverse a bad edge and from there, play adversarially against the primary strategy and co-operatively with an alternative strategy.

Claim 1. *If $\tau \in \mathfrak{S}_\forall$ is a winning strategy for Adam in \hat{G} , then there exist $\tau' \in \mathfrak{S}_\forall$ and $\sigma' \in \mathfrak{S}_\exists$ such that $\forall \sigma \in \mathfrak{S}_\exists : \mathbf{Val}(\sigma', \tau') - \mathbf{Val}(\sigma, \tau') \geq \lambda^{|V|} \min\{\mathbf{cVal}_{\neg v}^u(G) - w(u, v) - \lambda \mathbf{aVal}^v(G) : (u, v) \in \mathcal{B} \text{ and } u \in V_\exists\} > 0$.*

The claim follows from the definitions and Lemma 5. Conversely, winning strategies for Eve in \hat{G} are actually regret-free.

Claim 2. *If $\sigma \in \mathfrak{S}_\exists$ is a winning strategy for Eve in \hat{G} , then $\mathbf{reg}^\sigma(G) = 0$.*

Our argument to prove this claim requires we first show that a winning strategy for Eve ensures the antagonistic value of G from v_I . For completeness, a proof for this claim is included in appendix.

The desired result then follows from Lemma 1 and from the fact that membership of an edge in \mathcal{B} can be decided by computing \mathbf{cVal} and a threshold query regarding \mathbf{aVal} , thus in polynomial time. \square

We observe the proof of Theorem 2—more precisely, Claim 1—implies that, if there is no regret-free strategy for Eve in a game, then the regret of the game is at least $\lambda^{|V|}$ times the smallest local regret labelling the bad edge from \mathcal{B} which Adam can force. More formally:

Corollary 1. *If no regret-free strategy for Eve exists in G , then $\mathbf{Reg}(G) \geq a_G$ where $a_G := \lambda^{|V|} \min\{\mathbf{locreg}(uv, 0) : u \in V_\exists \text{ and } (u, v) \in \mathcal{B}\}$.*

Deciding r-regret. It will be useful in the sequel to define the *regret of a play* and the *regret of a play prefix*. Given a play $\pi = v_0 v_1 \dots$, we define the regret of π as:

$$\mathbf{reg}(\pi) := (\sup\{\mathbf{locreg}(\pi, i) : v_i \in V_\exists\} \cup \{0\}).$$

Intuitively, the local regrets give lower bounds for the overall regret of a play. We will also let the *regret of a play prefix* $\rho = v_0 \dots v_j$ be equal to

$$\max\left(\{\lambda^i (\mathbf{cVal}_{\neg v_{i+1}}^{v_i}(G) - \mathbf{Val}(\rho[i..j])) : 0 \leq i < j \text{ and } v_i \in V_\exists\} \cup \{0\}\right).$$

Let us give some more intuition regarding the regret of a play. Consider a pair of strategies σ and τ for Eve and Adam, respectively. Suppose there is an alternative strategy σ' for Eve, such that, against τ , the obtained payoff is greater than that of $\pi_{\sigma\tau}$. It should be clear that this implies there is some position i such that, from vertex $v_i \in V_\exists$ σ' and τ result in a different play from $\pi_{\sigma\tau}$ (see Figure 3). We will sometimes refer to this deviation, *i.e.* the play $\pi_{\sigma'\tau}$, as a *better alternative* to $\pi_{\sigma\tau}$.

We can now show the regret of a strategy for Eve in fact corresponds to the supremum of the regret of plays consistent with the strategy.

Lemma 6. *For any strategy σ of Eve, $\mathbf{reg}^\sigma(G) = \sup\{\mathbf{reg}(\pi) : \pi \text{ is consistent with } \sigma\}$.*

We note that for any play π , the sequence $\langle \lambda^i(\mathbf{cVal}_{\neg v_{i+1}}^{v_i}(G) - \mathbf{Val}(\pi[i..])) \rangle_{i \geq 0}$ converges to 0 because $(\mathbf{cVal}_{\neg v_{i+1}}^{v_i}(G) - \mathbf{Val}(\pi[i..]))$ is bounded by $\frac{2W}{(1-\lambda)}$. It follows that if we have a non-zero lower bound for the regret of π , then there is some index N such that the witness for the regret occurs before N . Moreover, we can place a polynomial upper bound on N . More precisely:

Lemma 7. *Let π be a play in G and suppose $0 < r \leq \mathbf{reg}(\pi)$. Let*

$$N(r) := \lfloor (\log r + \log(1 - \lambda) - \log(2W)) / \log \lambda \rfloor + 1.$$

Then $\mathbf{reg}(\pi) = \mathbf{reg}(\pi[..N(r)]) - \lambda^{N(r)} \mathbf{Val}(\pi[N(r)..])$.

The above result gives us a bound on how far we have to unfold a game after having witnessed a non-zero lower bound, r , for the regret. If we consider the example from Figure 3, this translates into a bound on how many turns after v_i a deviation can still yield bigger local regret (see Figure 4).

Corollary 1 then gives us the required lower bound to be able to use Lemma 7.

Lemma 8. *If $\mathbf{Reg}(G) \geq a_G$ then $\mathbf{Reg}(G)$ is equal to*

$$\inf_{\sigma \in \mathfrak{S}_\exists} \sup\{\mathbf{reg}(\pi[..N(a_G)]) - \lambda^{N(a_G)} \mathbf{aVal}^{v_{N(a_G)}}(G) : \pi = v_0 v_1 \dots \text{ is consistent with } \sigma\}.$$

This already implies we can compute the regret value in alternating polynomial time (or equivalently, deterministic polynomial space [5]).

Proposition 1. *The regret value is computable using only polynomial space.*

Proof. We first label the arena with the antagonistic and co-operative values and solve the safety game described for Theorem 2. The latter can be done in polynomial time. If the Eve wins the safety game, the regret value is 0. Otherwise, we know $a_G > 0$ is a lower bound for the regret value. We now simulate G using an alternating Turing machine which halts in at most $N(a_G)$ steps. That is, a polynomial number of steps. The simulated play prefix is then assigned a regret value as per Lemma 8 (recall we have already pre-computed the antagonistic value of every vertex). \square

As a side-product of the algorithm described in the above proof we get that finite memory strategies suffice for Eve to minimize her regret in a discounted-sum game.

Corollary 2. *Let $\mu := |\Delta|^{N(a_G)}$, with $N(0) = 0$. It holds that*

$$\mathbf{Reg}_{\Sigma_\exists^\mu, \mathfrak{S}_\forall}(G) = \mathbf{Reg}_{\mathfrak{S}_\exists, \mathfrak{S}_\forall}(G).$$

Simple regret-minimizing behaviours. We will now argue that Eve has a simple strategy which ensures regret of at most $\mathbf{Reg}(G)$. Her strategy will consist in “playing co-operatively” (*i.e.*, a strategy that attempts to maximize the co-operative payoff) for some turns (until a high local regret has already been witnessed) and then switch to a co-operative worst-case optimal strategy (*i.e.*, a strategy attempting to maximize the co-operative payoff while achieving at least the antagonistic payoff).

We will now define a family of strategies which switch from co-operative behaviour to antagonistic, after a specific number of turns have elapsed (in fact, enough for the discounted local regret to be less than the desired regret). Denote by σ^{co} a strongly co-operative strategy for Eve in G and by σ^{cw} a co-operative worst-case optimal strategy for Eve in G . Recall that, by Lemma 3, such strategies for her always exist. Finally, given a co-operative strategy σ^{co} , a co-operative worst-case optimal strategy σ^{cw} , and $t \in \mathbb{Q}$ let us define an *optimistic-then-pessimistic strategy for Eve* $[\sigma^{\text{co}} \xrightarrow{t} \sigma^{\text{cw}}]$. The strategy is such that, for any play prefix $\rho = v_0 \dots v_n$ such that $v_n \in V_\exists$

$$[\sigma^{\text{co}} \xrightarrow{t} \sigma^{\text{cw}}](\rho) = \begin{cases} \sigma^{\text{co}}(\rho) & \text{if } |\mathbf{cOpt}(v_n)| = 1 \text{ and } \mathbf{locreg}(\rho \cdot \sigma^{\text{cw}}(\rho), n+1) > t \\ \sigma^{\text{cw}}(\rho) & \text{otherwise.} \end{cases}$$

We claim that, when we set $t = \mathbf{Reg}(G)$, an optimistic-then-pessimistic strategy for Eve ensures minimal regret. That is

Proposition 2. *Let σ^{co} be a strongly co-operative strategy for Eve, σ^{cw} be a Eve and a co-operative worst-case optimal strategy for Eve, and $t = \mathbf{Reg}(G)$. The strategy $\sigma = [\sigma^{\text{co}} \xrightarrow{t} \sigma^{\text{cw}}]$ has the property that $\mathbf{reg}^\sigma(G) = \mathbf{Reg}(G)$.*

This is a refinement of the strategy one can obtain from applying the algorithm used to prove Proposition 1.¹ The latter tells us that a regret-minimizing strategy of Eve eventually switches to a worst-case optimal behaviour. For vertices where, before this switch, another edge was chosen by Eve, we argue that she must have been playing a co-operative strategy. Otherwise, she could have switched sooner. A full proof is provided in Appendix A.5.

We have shown the regret value can be computed using an algorithm which requires polynomial space only. This algorithm is based on a polynomial-length unfolding of the game and from it we can deduce that the regret value is representable using a polynomial number of bits. (Indeed, all exponents occurring in the formula from Lemma 8 will be polynomial according to Lemma 7.) Also, we have argued that Eve has a “simple” strategy σ to ensure minimal regret. Such a strategy is defined by two polynomial-time constructible sub-strategies and the regret value of the game. Hence, it can be encoded into a polynomial number of bits itself. Furthermore, σ is guaranteed to be playing as its co-operative worst-case optimal component after $N(\mathbf{Reg}(G))$ turns (see, again, Lemma 7), which is a polynomial number of turns. Given a regret threshold r , we claim we can verify that σ ensures regret at most r in polynomial time. This can be achieved by allowing Adam to play in G , and against σ , with the objective of reaching an edge with high local regret before $N(\mathbf{Reg}(G))$ turns. An possible formalization of this idea follows. Consider the product of G with a counter ranging from 1 to $N(\mathbf{Reg}(G))$ where we make all vertices belong to Adam. In this game H , we make edges leaving vertices previously belonging to Eve go to a sink and define a new weight function w' which assigns to these edges their negative non-discounted local regret: going from u to v when σ dictates to go to v' yields $w(u, v') + \lambda \mathbf{aVal}^{v'}(H \times \sigma) - w(u, v) + \lambda \mathbf{cVal}^v(H)$. Lemma 8 allows us to show that σ ensures regret at most r in G if and only if the antagonistic value of a discounted-sum game played on H with weight function w' is at most $-r$.

It follows that the regret threshold problem is in NP, as stated in Theorem 1.

Example 3. We revisit the discounted-sum game from Figure 1. Let us instantiate the values $M = 100$ and $\lambda = \frac{9}{10}$. According to our previous remarks on this arena, after i visits to v without Adam choosing (v, y) , Eve could achieve $(\frac{9}{10})^{2i}10$ by going to x or hope for $(\frac{9}{10})^{2i+1}1000$ by going to v again. Her best regret minimizing strategy corresponds to σ^{22} which ensures regret of at most $9.9030 = 10 - (\frac{9}{10})^{44}10$. It is easy to see that Eve cannot win the safety game \hat{G} constructed from this arena and that the lower bound a_G one can obtain from \hat{G} is equal to $1.2466 = (\frac{9}{10})^4(10 - (\frac{9}{10})^210)$. As expected, when Eve plays her optimal regret-minimizing (optimistic-then-pessimistic) strategy any better alternative must deviate before $N(a_G) = 71$ turns. Indeed, we have already argued that the regret 9.9030 is witnessed by Adam choosing the edge (v, y) for any strategy of Eve going to v more than 22 times.

4 Regret against positional strategies of Adam

In this section we consider the problem of computing the (minimal) regret when Adam is restricted to playing positional strategies.

Theorem 3. *Deciding if the regret value is less than a given threshold (strictly or non-strictly), playing against positional strategies of Adam, is in PSPACE.*

Playing against an Adam, when he is restricted to playing memoryless strategies gives Eve the opportunity to learn some of Adam’s strategic choices. However, due to its decaying nature, with the discounted-sum payoff function Eve must find a balance between exploring too quickly, thereby presenting lightly discounted alternatives; and learning too slowly, thereby heavily discounting her eventual payoff.

A similar approach to the one we have adopted in Section 3 can be used to obtain an algorithm for this setting. For reasons of space we defer its presentation to the appendix. The claimed lower bound follows from Theorem 5.

¹In fact, our proof of Prop. 2 relies in Eve requiring finite memory, to minimize her regret.

Deciding 0-regret. As in the previous section, we will reduce the problem of deciding if the game has regret value 0 to that of determining the winner of a safety game. It will be obvious that if no regret-free strategy for Eve exists in the original game, then we can construct, for any strategy of hers, a positional strategy of Adam which ensures non-zero regret. Hence, we will also obtain a lower bound on the regret of the game in the case Adam wins the safety game.

Let us fix some notation. For a set of edges $D \subseteq E$, we denote by $G|D$ the weighted arena $(V, V_\exists, v_I, D, w)$. Also, for a positional strategy $\tau : (V \setminus V_\exists) \rightarrow E$ for Adam in G , we denote by $G \times \tau$ the weighted arena resulting from removing all edges not consistent with τ . Next, for an edge $(s, t) \in E$ we define $E_\forall(st) := \{(u, v) \in E : \text{if } u = s \text{ then } v = t \text{ or } u \in V_\exists\}$. We extend the latter to play prefixes $\rho = v_0 \dots v_n$ by (recursively) defining $E_\forall(\rho) := E_\forall(\rho[..n-1]) \cap E_\forall(v_{n-1}v_n)$. If π is a play, then $E \supseteq E_\forall(\pi[..i]) \supseteq E_\forall(\pi[..j])$ for all $0 \leq i \leq j$. Hence, since E is finite, the value $E_\forall(\pi) := \lim_{i \geq 0} E_\forall(\pi[..i])$ is well-defined. Remark that $E_\forall(\pi)$ does not restrict edges leaving vertices of Eve. The following properties directly follow from our definitions.

Lemma 9. *Let π be a play or play prefix consistent with a positional strategy for Adam. It then holds that: (i) for every $v \in V \setminus V_\exists$ there is some edge $(v, \cdot) \in E_\forall(\pi)$, (ii) π is consistent with a strategy $\tau \in \Sigma_\forall^1(G)$ if and only if $\tau \in \Sigma_\forall^1(G|E_\forall(\pi))$, and (iii) every strategy $\tau \in \Sigma_\forall^1(G|E_\forall(\pi))$ is also an element from $\Sigma_\forall^1(G)$.*

To be able to decide whether regret-free strategies for Eve exist, we define a new safety game. The arena we consider is $\tilde{G} := (\hat{V}, \hat{V}_\exists, \hat{v}_I, \hat{E})$ where $\hat{V} := V \times \mathcal{P}(E)$, $\hat{V}_\exists := V_\exists \times \mathcal{P}(E)$, $\hat{v}_I := (v_I, E)$, and \hat{E} contains the edge $((u, C), (v, D))$ if and only if $(u, v) \in E$ and $D = C \cap E_\forall(uv)$.

Theorem 4. *Deciding if the regret value is 0, playing against positional strategies of Adam, is in PSPACE.*

Proof. A safety game is constructed as in the proof of Theorem 2. Here, we consider \tilde{G} and the set of bad edges $\tilde{\mathcal{B}} := \{((u, C), (v, D)) \in \hat{E} : u \in V_\exists \text{ and } \exists \tau \in \Sigma_\forall^1(G|C), w(u, v) + \lambda \mathbf{cVal}^v(G \times \tau) < \mathbf{cVal}_{-v}^u(G \times \tau)\}$. We then have the safety game $\tilde{G} = (\hat{V}, \hat{V}_\exists, \hat{v}_I, \hat{E}, \tilde{\mathcal{B}})$. Note that there is an obvious bijective mapping from plays (and play prefixes) in \tilde{G} to plays (prefixes) in G which are consistent with a positional strategy for Adam. One can then show the following properties hold:

Claim 3. *If $\tau \in \mathfrak{S}_\forall(\tilde{G})$ is a winning strategy for Adam in \tilde{G} , then for all $\sigma \in \mathfrak{S}_\exists(G)$, there exist $t_{\tau\sigma} \in \Sigma_\forall^1(G)$ and $s_{\tau\sigma} \in \mathfrak{S}_\exists(G)$ such that $\mathbf{Val}(s_{\tau\sigma}, t_{\tau\sigma}) - \mathbf{Val}(\sigma, t_{\tau\sigma}) \geq \lambda^{|V|(|E|+1)} \min\{\mathbf{cVal}_{-v}^u(G \times \tau) - w(u, v) - \lambda \mathbf{cVal}^v(G \times \tau) : ((u, C), (v, D)) \in \tilde{\mathcal{B}}, \tau \in \Sigma_\forall^1(G|C)\}$.*

The claim follows from *positional determinacy* of safety games and Lemma 9 (see Appendix B.1).

Claim 4. *If $\sigma \in \mathfrak{S}_\exists(\tilde{G})$ is a winning strategy for Eve in \tilde{G} , then there is $s_\sigma \in \mathfrak{S}_\exists(G)$ such that $\mathbf{reg}_{\mathfrak{S}_\exists, \Sigma_\forall^1}^{s_\sigma}(G) = 0$.*

It then follows from the determinacy of safety games that Eve wins the safety game \tilde{G} if and only if she has a regret-free strategy. We provide full proofs for these claims in appendix.

We observe that simple cycles in \tilde{G} have length at most $|V|(|E|+1)$. Thus, we can simulate the safety game until we complete a cycle and check that all traversed edges are good, all in alternating polynomial time. Indeed, an alternating Turing machine can simulate the cycle and then (universally) check that for all edges, for all positional strategies of the Adam, the inequality holds. \square

Corollary 3. *If no regret-free strategy for Eve exists in G , then $\mathbf{Reg}_{\mathfrak{S}_\exists, \Sigma_\forall^1}(G) \geq b_G$ where $b_G := \lambda^{|V|(|E|+1)} \min\{\mathbf{cVal}_{-v}^u(G \times \tau) - w(u, v) - \lambda \mathbf{cVal}^v(G \times \tau) : ((u, C), (v, D)) \in \tilde{\mathcal{B}} \text{ and } \tau \in \Sigma_\forall^1(G|C)\}$.*

Lower bounds. We claim that both 0-regret and r -regret are **coNP-hard**. This can be shown by adapting the reduction from 2-disjoint-paths given in [13] to the regret threshold problem against memoryless adversaries. For completeness, we provide the reductions here in appendix.

Theorem 5. *Let $\lambda \in (0, 1)$ and $r \in \mathbb{Q}$ be fixed. Deciding if the regret value is less than r (strictly or non-strictly), playing against positional strategies of Adam, is **coNP-hard**.*

5 Playing against word strategies of Adam

In this section, we consider the case where Adam is restricted to playing *word strategies*. First, we show that the regret threshold problem can be solved whenever the discounted sum automata associated to the game structure can be made deterministic. As the determinization problem for discounted sum automata has been solved in the literature for only sub-classes of discount factors, and left open in the general case, we complement this result by two other results. First, we show how to solve an ε -gap *promise* variant of the regret threshold problem, and second, we give an algorithm to solve the 0 regret problem. In the two cases, we obtain completeness results on the computational complexities of the problems.

Preliminaries. The formal definition of the ε -gap promise problem is given below. We first define here the necessary vocabulary. We say that a strategy of Adam is a *word strategy* if his strategy can be expressed as a function $\tau : \mathbb{N} \rightarrow [\max\{\mathbf{deg}^+(v) : v \in V\}]$, where $[n] = \{i : 1 \leq i \leq n\}$. Intuitively, we consider an order on the successors of each Adam vertex. On every turn, the strategy τ of Adam will tell him to move to the i -th successor of the vertex according to the fixed order. We denote by \mathfrak{W}_V the set of all such strategies for Adam. A game in which Adam plays word strategies can be reformulated as a game played on a weighted automaton $\Gamma = (Q, q_I, A, \Delta, w)$ and strategies of Adam—of the form $\tau : \mathbb{N} \rightarrow A$ —determine a sequence of input symbols, i.e. an omega word, to which Eve has to react by choosing Δ -successor states starting from q_I . In this setting a strategy of Eve which minimizes regret defines a run by resolving the non-determinism of Δ in Γ , and ensures the difference of value given by the constructed run is minimal w.r.t. to the value of the best run on the word spelled out by Adam.

Deciding 0-regret. We will now show that if the regret of an arena (or automaton) is 0, then we can construct a memoryless strategy for Eve which ensures no regret is incurred. More specifically, assuming the regret is 0, we have the existence of a family of strategies of Eve which ensure decreasing regret (with limit 0). We use this fact to choose a small enough ε and the corresponding strategy of hers from the aforementioned family to construct a memoryless strategy for Eve with nice properties which allow us to conclude that its regret is 0. Hence, it follows that an automaton has zero regret if and only if a memoryless strategy of Eve ensures regret 0. As we can guess such a strategy and easily check if it is indeed regret-free (using the obvious reduction to non-emptiness of discounted-sum automata or one-player discounted-sum games), the problem is in NP. A matching lower bound follows from a reduction from SAT which was first described in [1]. We sketch it, for completeness, in the appendix.

Theorem 6. *Deciding if the regret value is 0, playing against word strategies of Adam, is NP-complete.*

Deciding r-regret: determinizable cases. When the weighted automaton Γ associated to the game structure can be made deterministic, we can solve the regret threshold problem with the following algorithm. In [13] we established that, against eloquent adversaries, computing the regret reduced to computing the value of a quantitative simulation game as defined in [6]. The game is obtained by taking the product of the original automaton and a deterministic version of it. The new weight function is the difference of the weights of both components (for each pair of transitions). In [4], it is shown how to determinize discounted-sum automata when the discount factor is of the form $\frac{1}{n}$, for $n \in \mathbb{N}$. So, for this class of discount factor, we can state the following theorem:

Theorem 7. *Deciding if the regret value is less than a given threshold (strictly or non-strictly), playing against word strategies of Adam, is in EXPTIME for λ of the form $\frac{1}{n}$.*

The ε -gap promise problem. Given a discounted-sum automaton \mathcal{A} , $r \in \mathbb{Q}$, and $\varepsilon > 0$, the ε -gap promise problem adds to the regret threshold problem the hypothesis that \mathcal{A} will either have regret $\leq r$ or $> r + \varepsilon$. We observe that an algorithm which gives:

- a YES answer implies that $\mathbf{Reg}_{\Sigma_\exists, \mathfrak{W}_V}(\mathcal{A}) \leq r + \varepsilon$,
- whereas a NO answer implies $\mathbf{Reg}_{\Sigma_\exists, \mathfrak{W}_V}(\mathcal{A}) > r$.

will decide the ε -gap promise problem.

In [4], it is shown that there are discounted-sum automata which define functions that cannot be realized with deterministic-sum automata. Nevertheless, it is also shown in that paper that given a discounted-sum automaton it is always possible to construct a deterministic one that is ε -close in the following formal sense. A *discounted-sum automaton* \mathcal{A} is ε -close to another discounted sum automaton \mathcal{B} , if for all words x the absolute value of the difference between the values assign by \mathcal{A} and \mathcal{B} to x is at most ε . So, it should be clear that we can apply the algorithm underlying Theorem 7 to Γ and a determinized version \mathcal{D}_Γ of it (which is ε -close to Γ) and solve the ε -gap promise problem. We can then prove the following result.

Theorem 8. *Deciding the ε -gap regret problem is in PSPACE.*

The complexity of the algorithm follows from the fact that the value of a (quantitative simulation) game, played on the product of Γ and \mathcal{D}_Γ we described above, can be determined by simulating the game for a polynomial number of turns. Thus, although the automaton constructed using the techniques of Boker and Henzinger [4] is of size exponential, we can construct it “on-the-fly” for the required number of steps and then stop.

Lower bounds. We claim the ε -gap promise problem is PSPACE-hard even if both λ and ε are not part of the input. To establish the result, we give a reduction from QSAT which uses the gadgets depicted in Figures 11 and 12. For space reasons we defer the reduction to Appendix C.

Theorem 9. *Let $\lambda \in (0, 1)$ and $\varepsilon \in (0, 1)$ be fixed. As input, assume we are given $r \in \mathbb{Q}$ and weighted arena \mathcal{A} such that $\mathbf{Reg}_{\Sigma_\exists, \mathbb{W}_\forall}(\mathcal{A}) \leq r$ or $\mathbf{Reg}_{\Sigma_\exists, \mathbb{W}_\forall}(\mathcal{A}) > r + \varepsilon$. Deciding if the regret value is less than a given threshold, playing against word strategies of Adam, is PSPACE-hard.*

It follows that the general problem is also PSPACE-hard (even if ε is set to 0).

Corollary 4. *Let $\lambda \in (0, 1)$. For $r \in \mathbb{Q}$, weighted arena G , determining whether $\mathbf{Reg}_{\Sigma_\exists, \mathbb{W}_\forall}(G) \triangleleft r$, for $\triangleleft \in \{<, \leq\}$, is PSPACE-hard.*

References

- [1] Benjamin Aminof, Orna Kupferman, and Robby Lampert. Reasoning about online algorithms with weighted automata. *ACM Transactions on Algorithms*, 2010.
- [2] Krzysztof R. Apt and Erich Grädel. *Lectures in game theory for computer scientists*. Cambridge University Press, 2011.
- [3] David E. Bell. Regret in decision making under uncertainty. *Operations Research*, 30(5):961–981, 1982.
- [4] Udi Boker and Thomas A. Henzinger. Exact and approximate determinization of discounted-sum automata. *LMCS*, 10(1), 2014.
- [5] Ashok K. Chandra, Dexter Kozen, and Larry J. Stockmeyer. Alternation. *J. ACM*, 28(1):114–133, 1981.
- [6] Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Quantitative languages. *ACM Transactions on Computational Logic*, 11(4), 2010.
- [7] Werner Damm and Bernd Finkbeiner. Does it pay to extend the perimeter of a world model? In *FM*, volume 6664 of *LNCs*, pages 12–26. Springer, 2011.
- [8] Tali Eilam-Tzoref. The disjoint shortest paths problem. *Discrete Applied Mathematics*, 85(2):113–138, 1998.
- [9] Emmanuel Filiot, Tristan Le Gall, and Jean-François Raskin. Iterated regret minimization in game graphs. In *MFCS*, volume 6281 of *LNCs*, pages 342–354. Springer, 2010.
- [10] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.

- [11] Thomas A. Henzinger and Nir Piterman. Solving games without determinization. In *CSL*, pages 395–410, 2006.
- [12] Paul Hunter, Guillermo A. Pérez, and Jean-François Raskin. Mean-payoff games with partial-observation - (extended abstract). In *RP*, pages 163–175, 2014.
- [13] Paul Hunter, Guillermo A. Pérez, and Jean-François Raskin. Reactive synthesis without regret. In Luca Aceto and David de Frutos-Escrig, editors, *CONCUR*, volume 42 of *LIPIcs*, pages 114–127. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.
- [14] Daniel Dominic Sleator and Robert Endre Tarjan. Amortized efficiency of list update rules. In *Proceedings of the 16th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1984, Washington, DC, USA*, pages 488–492. ACM, 1984.
- [15] Uri Zwick and Mike Paterson. The complexity of mean payoff games on graphs. *TCS*, 158(1):343–359, 1996.

A Missing Proofs From Section 3

A.1 Proof of Lemma 6

Consider any $\sigma, \sigma' \in \mathfrak{S}_\exists$ and $\tau \in \mathfrak{S}_\forall$ such that $\pi_{\sigma\tau} \neq \pi_{\sigma'\tau}$. Let us write $\pi_{\sigma\tau} = v_0v_1\dots$ and $\pi_{\sigma'\tau} = v'_0v'_1\dots$ and denote by ℓ the length of the longest common prefix of $\pi_{\sigma\tau}$ and $\pi_{\sigma'\tau}$. We claim that

$$\lambda^\ell(\mathbf{cVal}_{\neg v_{\ell+1}}^{v_\ell}(G) - \mathbf{Val}(\pi_{\sigma\tau})) \geq \lambda^\ell(\mathbf{Val}(\pi_{\sigma'\tau}[\ell..]) - \mathbf{Val}(\pi_{\sigma\tau}[\ell..])). \quad (1)$$

Indeed, if we assume it is not the case, we then get that $\mathbf{cVal}_{\neg v_{\ell+1}}^{v_\ell}(G) < \mathbf{Val}(\pi_{\sigma'\tau}[\ell + 1..])$, which contradicts the definition of \mathbf{cVal} . Note that Lemma 4 actually tells us that there is another strategy τ' for Adam and a second alternative strategy σ'' for Eve which give us equality in the above equation. More formally, from Equation 1 and Lemma 4 we get that for all $\sigma \in \mathfrak{S}_\exists$, if there are $\tau \in \mathfrak{S}_\forall$ and $\sigma' \in \mathfrak{S}_\exists$ such that $\pi_{\sigma\tau} \neq \pi_{\sigma'\tau}$ then

$$\sup_{\tau, \sigma' \text{ s.t. } \pi_{\sigma\tau} \neq \pi_{\sigma'\tau}} \lambda^\ell(\mathbf{Val}(\pi_{\sigma'\tau}[\ell..]) - \mathbf{Val}(\pi_{\sigma\tau}[\ell..])) = \lambda^\ell(\mathbf{cVal}_{\neg v_{\ell+1}}^{v_\ell}(G) - \mathbf{Val}(\pi_{\sigma\tau})). \quad (2)$$

We are now able to prove the result. That is, for any strategy σ for Eve:

$$\begin{aligned} & \sup\{\mathbf{reg}(\pi) : \pi \text{ is consistent with } \sigma\} \\ &= \sup_{\tau \in \mathfrak{S}_\forall} \mathbf{reg}(\pi_{\sigma\tau} = v_0v_1\dots) \quad \text{def. of } \pi_{\sigma\tau} \\ &= \sup_{\tau \in \mathfrak{S}_\forall} \max \left\{ 0, \sup_{\substack{i \geq 0 \\ v_i \in V_\exists}} \lambda^i(\mathbf{cVal}_{\neg v_{i+1}}^{v_i}(G) - \mathbf{Val}(\pi_{\sigma\tau}[i..])) \right\} \quad \text{def. of } \mathbf{reg}(\pi_{\sigma\tau}) \\ &= \sup_{\tau \in \mathfrak{S}_\forall} \max \left\{ 0, \sup_{\sigma' \text{ s.t. } \pi_{\sigma\tau} \neq \pi_{\sigma'\tau}} \lambda^\ell(\mathbf{Val}(\pi_{\sigma'\tau}[\ell..]) - \mathbf{Val}(\pi_{\sigma\tau}[\ell..])) \right\} \quad \text{by Eq. (2)} \\ &= \sup_{\tau \in \mathfrak{S}_\forall} \max \left\{ 0, \sup_{\sigma' \text{ s.t. } \pi_{\sigma\tau} \neq \pi_{\sigma'\tau}} (\mathbf{Val}(\sigma', \tau) - \mathbf{Val}(\sigma, \tau)) \right\} \quad \text{def. of } \mathbf{Val}(\cdot), \ell \\ &= \sup_{\tau \in \mathfrak{S}_\forall} \sup_{\sigma' \in \mathfrak{S}_\exists} (\mathbf{Val}(\sigma', \tau) - \mathbf{Val}(\sigma, \tau)) \quad 0 \text{ when } \pi_{\sigma\tau} = \pi_{\sigma'\tau} \end{aligned}$$

as required. \square

A.2 Proof of Lemma 7

Observe that $N(r)$ is such that $\frac{2W\lambda^{N(r)}}{1-\lambda} < r$. Hence, we have that for all $i \geq N(r)$ such that $v_i \in V_\exists$ it holds that $\lambda^i(\mathbf{cVal}_{\neg v_{i+1}}^{v_i}(G) - \mathbf{Val}(\pi[i..])) \leq \frac{2W\lambda^{N(r)}}{1-\lambda} < r$. It follows that

$$\begin{aligned} \mathbf{reg}(\pi) &= \sup\{\lambda^i(\mathbf{cVal}_{\neg v_{i+1}}^{v_i}(G) - \mathbf{Val}(\pi[i..])) : i \geq 0 \text{ and } v_i \in V_\exists\} \\ &= \max_{\substack{0 \leq i < N(r) \\ v_i \in V_\exists}} \lambda^i(\mathbf{cVal}_{\neg v_{i+1}}^{v_i}(G) - \mathbf{Val}(\pi[i..N(r)])) - \lambda^{N(r)}\mathbf{Val}(\pi[N(r)..]) \end{aligned}$$

as required. \square

A.3 Proof of Lemma 8

First, note that if $\mathbf{Reg}(G) > 0$ then there cannot be any regret-free strategies for Eve in G . It then follows from Corollary 1 that $\mathbf{Reg}(G) \geq a_G$. Next, using Lemma 7 and the definition of the regret of a play we have that $\mathbf{Reg}(G)$ is equal to

$$\inf_{\sigma \in \mathfrak{S}_\exists} \sup\{\mathbf{reg}(\pi[..N(a_G)]) - \lambda^{N(a_G)}\mathbf{Val}(\pi[N(a_G)..]) : \pi \text{ is consistent with } \sigma\}.$$

Finally, note that it is in the interest of Eve to maximize the value $\lambda^{N(a_G)} \mathbf{Val}(\pi[N(a_G)..])$ in order to minimize regret. Conversely, Adam tries to minimize the same value. Thus, we can replace it by the antagonistic value from $\pi[N(a_G)..]$ discounted accordingly. More formally, we have

$$\begin{aligned}
& \inf_{\sigma \in \mathfrak{S}_\exists} \sup \{ \mathbf{reg}(\pi[..N(a_G)]) - \lambda^{N(a_G)} \mathbf{Val}(\pi[N(a_G)..]) : \pi \text{ is consistent with } \sigma \} \\
&= \inf_{\sigma \in \mathfrak{S}_\exists} \sup_{\tau \in \mathfrak{S}_\forall} \mathbf{reg}(\pi_{\sigma\tau}[..N(a_G)]) - \lambda^{N(a_G)} \mathbf{Val}(\pi_{\sigma\tau}[N(a_G)..]) \\
&= \inf_{\substack{\sigma \in \mathfrak{S}_\exists \\ \sigma' \in \mathfrak{S}_\exists}} \sup_{\substack{\tau \in \mathfrak{S}_\forall \\ \tau' \in \mathfrak{S}_\forall}} \mathbf{reg}(\pi_{\sigma\tau}[..N(a_G)] = \dots v) - \lambda^{N(a_G)} \mathbf{Val}^v(\sigma', \tau') \\
&= \inf_{\sigma \in \mathfrak{S}_\exists} \sup_{\tau \in \mathfrak{S}_\forall} \mathbf{reg}(\pi_{\sigma\tau}[..N(a_G)] = \dots v) + \inf_{\sigma' \in \mathfrak{S}_\exists} \sup_{\tau' \in \mathfrak{S}_\forall} \left(-\lambda^{N(a_G)} \mathbf{Val}^v(\sigma', \tau') \right) \\
&= \inf_{\sigma \in \mathfrak{S}_\exists} \sup_{\tau \in \mathfrak{S}_\forall} \mathbf{reg}(\pi_{\sigma\tau}[..N(a_G)] = \dots v) - \lambda^{N(a_G)} \left(\sup_{\sigma' \in \mathfrak{S}_\exists} \inf_{\tau' \in \mathfrak{S}_\forall} \mathbf{Val}^v(\sigma', \tau') \right) \\
&= \inf_{\sigma \in \mathfrak{S}_\exists} \sup_{\tau \in \mathfrak{S}_\forall} \mathbf{reg}(\pi_{\sigma\tau}[..N(a_G)] = \dots v) - \lambda^{N(a_G)} \mathbf{aVal}^v(G)
\end{aligned}$$

as required. \square

A.4 Proof of Claim 2

As a first step towards proving the result, we first make the observation that any winning strategy of Eve in \hat{G} also ensures a value of at least $\mathbf{aVal}(G)$ in the discounted-sum game played on G . More formally,

Claim 5. *If $\sigma \in \mathfrak{S}_\exists$ is a winning strategy for Eve in \hat{G} , then*

$$\forall \tau \in \mathfrak{S}_\forall, \forall i \geq 0 : \mathbf{Val}(\pi_{\sigma\tau}[i..] = v_i \dots) \geq \mathbf{aVal}^{v_i}(G). \quad (3)$$

Proof. Consider a winning strategy $\sigma \in \mathfrak{S}_\exists$ for Eve in \hat{G} . Since safety games are positionally determined (see, *e.g.* [2]) we can assume w.l.o.g. that σ is memoryless.

To convince the reader that σ has the property from Equation (3), we consider the synchronized product of G and σ —that is, the synchronized product of G and the finite Moore machine realizing σ . As σ is memoryless, then this product, which we denote in the sequel by $G \times \sigma$, is finite. Now, towards a contradiction, suppose that Equation (3) does not hold for σ . Further, let us consider an alternative (memoryless) strategy σ' of Eve which ensures $\mathbf{aVal}^v(G)$ from all $v \in V$. The latter exists by definition of $\mathbf{aVal}(G)$ and memoryless determinacy of discounted-sum games (see, *e.g.* [15]).

Let H denote a copy of $G \times \sigma$ where all edges induced by E are added—not just the ones allowed by σ —and $H|_{\sigma'}$ denote the sub-graph of H where only edges allowed by σ' are left. Since, by assumption, σ does not have the property of Equation (3) then the edges present in at least one vertex from $H|_{\sigma'}$ and $G \times \sigma$ differ. Note that such a vertex u is necessarily such that $u \in V_\exists$. Furthermore, from our definition of a strategy, we know that there is a single outgoing edge from it in both structures. Let us write (u, v) for the edge in $G \times \sigma$ and (u, v') for the edge in $H|_{\sigma'}$. Recall that σ is winning for Eve in \hat{G} . Thus, we have that $(u, v) \notin \mathcal{B} = \{(u, v) \in E : u \in V_\exists \text{ and } w(u, v) + \lambda \mathbf{aVal}^v(G) < \mathbf{cVal}_{\neg v}^u(G)\}$. It follows that

$$\begin{aligned}
w(u, v) + \lambda \mathbf{aVal}^v(H) &\geq \max_{x \neq v} \{w(u, x) + \lambda \mathbf{cVal}^x(H)\} \\
&\geq \max_{x \neq v} \{w(u, x) + \lambda \mathbf{aVal}^x(H)\} && \text{as } \mathbf{cVal}^x(H) \geq \mathbf{aVal}^x(H) \\
&= \mathbf{aVal}^u(H) && \text{because } u \in V_\exists.
\end{aligned}$$

Thus, the strategy σ'' of Eve which takes (u, v) instead of (u, v') and follows σ' otherwise—indeed, this might mean σ'' is not memoryless—also achieves at least $\mathbf{aVal}^u(H)$ from u onwards and is therefore an worst-case optimal antagonistic strategy in G (*i.e.* it has the property of Equation (3)). Notice that this process can be repeated for all vertices in which the two structures differ. Further, since both are finite, it will eventually terminate and yield a strategy of Eve which plays exactly as σ and for which Equation (3) holds, which is absurd. \square

Once more, consider a winning strategy $\sigma \in \mathfrak{S}_\exists$ for Eve in \hat{G} . We will now show that

$$\forall \tau \in \mathfrak{S}_\forall, \forall \sigma' \in \mathfrak{S}_\exists \setminus \{\sigma\} : \mathbf{Val}(\sigma, \tau) \geq \mathbf{Val}(\sigma', \tau).$$

The desired result will then directly follow.

Consider arbitrary strategies $\tau \in \mathfrak{S}_\forall$ and $\sigma' \in \mathfrak{S}_\exists \setminus \{\sigma\}$. Suppose that $\pi_{\sigma\tau} \neq \pi_{\sigma'\tau}$, as our claim trivially holds otherwise. Let ι be the maximal index $i \geq 0$ such that, if we write $\pi_{\sigma\tau} = v_0 v_1 \dots$ and $\pi_{\sigma'\tau} = v'_0 v'_1 \dots$, then $v_i = v'_i$. That is, ι is the maximal index for which the outcomes of σ and τ , and σ' and τ coincide. Note that v_ι is necessarily an Eve vertex, *i.e.* $v_\iota \in V_\exists$. We observe that, by definition of \mathbf{cVal} , it holds that

$$\mathbf{Val}(\pi_{\sigma'\tau}[\iota + 1..]) \leq \mathbf{cVal}^{v'_{\iota+1}}(G). \quad (4)$$

Furthermore, we know from the fact that σ is winning for Eve in \hat{G} that the edge $(v_\iota, v_{\iota+1})$ is such that

$$w(v_\iota, v_{\iota+1}) + \lambda \mathbf{aVal}^{v_{\iota+1}}(G) \geq \max_{t \neq v_{\iota+1}} \{w(v_\iota, t) + \lambda \mathbf{cVal}^t(G)\}. \quad (5)$$

In particular, this implies that $w(v_\iota, v_{\iota+1}) + \lambda \mathbf{aVal}^{v_{\iota+1}}(G) \geq w(v_\iota, v'_{\iota+1}) + \lambda \mathbf{cVal}^{v'_{\iota+1}}(G)$. It is then easy to verify that $w(v_\iota, v_{\iota+1}) + \lambda \mathbf{aVal}^{v_{\iota+1}}(G) = \mathbf{aVal}^{v_\iota}(G)$ using the observation that $v_\iota \in V_\exists$. From Claim 5 we also get that

$$\mathbf{Val}(\pi_{\sigma\tau}[\iota..]) \geq \mathbf{aVal}^{v_\iota}(G). \quad (6)$$

Putting all the above inequalities together, we have

$$\begin{aligned} \mathbf{Val}(\pi_{\sigma\tau}[\iota..]) &\geq \mathbf{aVal}^{v_\iota}(G) = w(v_\iota, v_{\iota+1}) + \lambda \mathbf{aVal}^{v_{\iota+1}}(G) && \text{by Eqn. (6)} \\ &\geq w(v_\iota, v'_{\iota+1}) + \lambda \mathbf{cVal}^{v'_{\iota+1}}(G) && \text{by Eqn. (5)} \\ &\geq \mathbf{Val}(\pi_{\sigma'\tau}[\iota..]) && \text{by Eqn. (4)} \end{aligned}$$

which, in turn, implies $\mathbf{Val}(\sigma, \tau) \geq \mathbf{Val}(\sigma', \tau)$ since $\pi_{\sigma\tau}[\iota..] = \pi_{\sigma'\tau}[\iota..]$. \square

A.5 Proof of Proposition 2

Let us start by showing that the regret of a play π is bounded (from above) by the discounted local regret from any index i , where from the i -th turn onwards Eve plays a worst-case optimal strategy. More formally:

Lemma 10. *Let $\pi = v_0 v_1 \dots$ be a play. Assume there is some $i \in \mathbb{N}$ such that*

- (i) $v_i \in V_\exists$;
- (ii) $\mathbf{reg}(\pi) \leq \lambda^i \mathbf{reg}(\pi[i..])$; and
- (iii) $\mathbf{aVal}^{v_j}(G) = w(v_j, v_{j+1}) + \lambda \mathbf{aVal}^{v_{j+1}}(G)$, for all $j \geq i$.

It then holds that $\mathbf{reg}(\pi) \leq \lambda^i (\mathbf{cVal}^{v_i}(G) - \mathbf{aVal}^{v_i}(G))$.

Proof. If $\mathbf{reg}(\pi) = 0$ then the claim holds trivially. Hence, let us assume $\mathbf{reg}(\pi) > 0$. It follows from Lemma 8 and Assumption (ii) that there exists $k \geq i$ such that $v_k \in V_\exists$ and

$$\mathbf{reg}(\pi) = \lambda^k (\mathbf{cVal}^{v_k}(G) - w(v_k, v_{k+1}) - \lambda \mathbf{aVal}^{v_{k+1}}(G)).$$

Observe that $\mathbf{cVal}^{v_k}(G) \geq \mathbf{cVal}^{v_k}_{\neg v_{k+1}}(G)$, by definition, and that from Assumption (iii) we have that $\mathbf{aVal}^{v_k}(G) \leq w(v_k, v_{k+1}) + \lambda \mathbf{aVal}^{v_{k+1}}(G)$. Thus, we get that $\mathbf{reg}(\pi) \leq \lambda^k (\mathbf{cVal}^{v_k}(G) - \mathbf{aVal}^{v_k}(G))$. Also, note that by definition of \mathbf{cVal} we have that

$$\mathbf{cVal}^{v_j}(G) \geq w(v_j, v_{j+1}) + \lambda \mathbf{cVal}^{v_{j+1}}(G)$$

for all $j \geq 0$. It thus follows from Assumption (iii) and the previous arguments that $\mathbf{reg}(\pi) \leq \lambda^i (\mathbf{cVal}^{v_i}(G) - \mathbf{aVal}^{v_i}(G))$ as required. \square

We are now ready to prove the Proposition holds.




Figure 5: Gadget to reduce a game to its regret game.

The zero case. If $\mathbf{Reg}(G) = 0$, then it follows from our reduction to safety games that Eve has a co-operative worst-case optimal strategy which minimizes regret. Indeed, it is straightforward to show that the strategy for Eve obtained from the safety game does not only ensure at least the antagonistic value, but it is also co-operative worst-case optimal. Thus, since $[\sigma^{\text{co}} \xrightarrow{0} \sigma^{\text{cw}}]$ is clearly equivalent to σ^{cw} in this case, the result follows.

Non-zero regret. Let us assume that $\mathbf{Reg}(G) > 0$. It then follows from Lemma 8 that Eve has a finite memory strategy σ which ensures regret of at most $\mathbf{Reg}(G)$ (see Corollary 2) and which, furthermore, can be assumed to switch after turn $N(a_G)$ to a co-operative worst-case optimal strategy σ^{cw} for Eve (since such a strategy ensures at least the antagonistic value of the vertex from which Eve starts playing it). We will further assume, w.l.o.g., that for all play prefixes $\pi = v_0 \dots v_n$ with $n \leq N(a_G)$, $v_n \in V_{\exists}$ and having $\sigma^{\text{cw}}(\pi) \neq \sigma^{\text{co}}(\pi) = \sigma(\pi)$, if σ switches to σ^{cw} from π onwards—that is, for all prefixes extending π —then the regret of the resulting strategy is strictly greater than $\mathbf{Reg}(G)$. Otherwise, one can consider the strategy resulting from the previously described switch instead of σ .

We will now argue that for all play prefixes $\pi = v_0 \dots v_n$ with $n \leq N(a_G)$ and $v_n \in V_{\exists}$, if $\sigma(\pi) \neq \sigma^{\text{cw}}$ then $\mathbf{cOpt}(v_n)$ is a singleton and $\mathbf{locreg}\pi[..n] \cdot \sigma^{\text{cw}}(\pi[..n])n + 1 > \mathbf{Reg}(G)$. The desired result will follow since in order for our assumption of $\mathbf{reg}(\sigma) = \mathbf{Reg}(G)$ to be true Eve must then choose the unique edge leading to the single element in $\mathbf{cOpt}(v_n)$.

Let us consider two cases.

First, if $\mathbf{locreg}\pi[..n] \cdot \sigma^{\text{cw}}(\pi[..n])n + 1 \leq \mathbf{Reg}(G)$, we can switch to σ^{cw} from $\pi[..n]$ onwards. Contradicting our initial assumption.

Second, if $|\mathbf{cOpt}(v_n)| > 1$ and $\mathbf{locreg}\pi[..n] \cdot \sigma^{\text{cw}}(\pi[..n])n + 1 > \mathbf{Reg}(G)$, then by Lemma 10 we get that the regret of the play (if we switched to σ^{cw}) is bounded above by $\lambda^n (\mathbf{cVal}^{v_n}(G) - \mathbf{aVal}^{v_n}(G))$. Also, since $\mathbf{cOpt}(v_n)$ is not a singleton, if Eve does not switch, then she cannot ensure a local regret of less than $\lambda^n (\mathbf{cVal}^{v_n}(G) - \mathbf{aVal}^{v_n}(G))$ —particularly, not even by taking an edge leading to a vertex in $\mathbf{cOpt}(v_n)$. This contradicts the assumption that that switching to σ^{cw} yields strictly more regret. \square

A.6 Lower bound

We now establish a lower bound for computing the minimal regret against any strategy by reducing from the problem of determining the antagonistic value of a discounted-sum game. More precisely, from a weighted arena G we construct, in logarithmic space, a weighted arena G' such that the antagonistic value of G is equal to the regret value of G' . This gives us:

Lemma 11. *Computing the regret of a discounted-sum game is at least as hard as computing the antagonistic value of a (polynomial-size) game with the same payoff function.*

Proof of Lemma 11. Suppose G is a weighted arena with initial vertex v_I . Consider the weighted arena G' obtained by adding to G the gadget of Figure 5 with $K := \frac{W}{1-\lambda}$. The initial vertex of G' is set to be v'_I . We will show that $\mathbf{aVal}(G) = K + 1 - \mathbf{Reg}(G')/\lambda$.

At v'_I Eve has a choice: she can choose to remain in the gadget or she can move to the original game G . If Eve remains in the gadget her payoff will be $\lambda(-3K - 2)$ while Adam could choose to enter the game and achieve a payoff of $\lambda \cdot \mathbf{cVal}(G)$. In this case her regret is $\lambda(\mathbf{cVal}(G) + 3K + 2) \geq \lambda(2K + 2)$.

Otherwise, if she chooses to play into G she can achieve at most $\lambda \cdot \mathbf{aVal}(G)$. The strategy of Adam which maximizes regret against this choice of Eve is the one which remains in the gadget. The payoff for Adam is $\lambda(K+1)$ in this case. Hence, the regret of the game in this scenario is $\lambda(K+1 - \mathbf{aVal}(G)) \leq \lambda(2K+1)$. Clearly she will choose to enter the game and $\mathbf{Reg}(G') = \lambda(K+1 - \mathbf{aVal}(G))$. \square

B Missing Proofs from Section 4

B.1 Proof of Claim 3

We will now argue that if $\tau \in \mathfrak{S}_V(\tilde{G})$ is a winning strategy for Adam in \tilde{G} , then for all $\sigma \in \mathfrak{S}_\exists(G)$, there exist $t_{\tau\sigma} \in \Sigma_V^1(G)$ and $s_{\tau\sigma} \in \mathfrak{S}_\exists(G)$ such that $\mathbf{Val}(s_{\tau\sigma}, t_{\tau\sigma}) - \mathbf{Val}(\sigma, t_{\tau\sigma})$ is at least

$$\lambda^{|V|(|E|+1)} \min_{\substack{((u,C),(v,D)) \in \tilde{\mathcal{B}} \\ \tau \in \Sigma_V^1(G|C)}} \{ \mathbf{cVal}_{\neg v}^u(G \times \tau) - w(u,v) - \lambda \mathbf{cVal}^v(G \times \tau) \}. \quad (7)$$

The argument is straightforward and based on the bijection between plays from G , which are consistent with positional strategies of Adam, and plays in \tilde{G} . Recall that safety games are positionally determined. That is, either Eve has a positional strategy which allows her to perpetually avoid the unsafe edges against any strategy for Adam, or Adam has a positional strategy which ensures that—regardless of the behaviour of Eve—the play eventually traverses some unsafe edge. Thus, since we assume $\tau \in \mathfrak{S}_V(\tilde{G})$ is winning for Adam in \tilde{G} we can assume that τ is in fact a positional strategy for Adam in \tilde{G} . Now consider an arbitrary strategy σ for Eve in G . We note, once more, that τ is a strategy for Adam in G , not only in \tilde{G} . Furthermore, τ is a positional strategy for Adam in G . Conversely, σ is a valid strategy for Eve in \tilde{G} . These facts follow from the definition of $E_V(\cdot)$ and construction \tilde{G} . Since τ is winning for Adam in \tilde{G} , the play $\pi_{\sigma\tau}$ traverses an unsafe edge. In fact, since τ is positional, the unsafe edge is necessarily traversed in at most $|V|(|E|+1)$ steps—that is, at most the length of the longest simple path in \tilde{G} . Let us write $(\tilde{v}_i, \tilde{v}_{i+1}) = ((v_i, C_i), (v_{i+1}, C_{i+1}))$ for the traversed unsafe edge at step $i \leq |V|(|E|+1)$. By definition of $\tilde{\mathcal{B}}$ we have that there exists $t_{\tau\sigma} \in \Sigma_V^1(G|C_i)$ such that

$$\mathbf{cVal}_{\neg v_{i+1}}^{v_i}(G \times t_{\tau\sigma}) - w(v_i, v_{i+1}) - \lambda \mathbf{cVal}^{v_i}(G \times t_{\tau\sigma}).$$

We now move from the game \tilde{G} back to the original game G . Henceforth, we consider the play $\pi_{\sigma\tau} = v_0v_1\dots$ in G which corresponds to $\pi_{\sigma\tau} = (v_0, C_0)(v_1, C_1)\dots$ in \tilde{G} . It is easy to see that $\pi_{\sigma\tau}[\dots i]$ is consistent with $t_{\tau\sigma}$. Hence, $\pi_{\sigma t_{\tau\sigma}}$ traverses edge (v_i, v_{i+1}) corresponding to bad edge $(\tilde{v}_i, \tilde{v}_{i+1})$ in \tilde{G} . Finally, by determinacy of discounted-sum games and by virtue of $G \times t_{\tau\sigma}$ being a finite weighted arena, we have that there is a strategy $s_{\tau\sigma} \in \mathfrak{S}_\exists(G \times t_{\tau\sigma})$ such that $\mathbf{Val}_G^{v_i}(s_{\tau\sigma}, t_{\tau\sigma}) = \mathbf{cVal}^{v_i}(G \times t_{\tau\sigma})$. It then follows from the definition of \mathbf{cVal} and $G \times s_{\tau\sigma}$ that $\mathbf{Val}_G^{v_i}(s_{\tau\sigma}, t_{\tau\sigma}) - \mathbf{Val}_G^{v_i}(\sigma, t_{\tau\sigma})$ is at least the value from Equation (7), just as required. \square

B.2 Proof of Claim 4

Let us show that if $\sigma \in \mathfrak{S}_\exists(\tilde{G})$ is a winning strategy for Eve in \tilde{G} , then there is $s_\sigma \in \mathfrak{S}_\exists(G)$ such that $\mathbf{reg}_{\mathfrak{S}_\exists, \Sigma_V^1}^{s_\sigma}(G) = 0$. The intuition behind the argument is the same as for the proof of Claim 2. However, in this case we first need to describe how to construct the strategy for Eve in G from a strategy for her in \tilde{G} .

A regret-free strategy from \tilde{G} . Observe that, by construction of \tilde{G} , for any vertex $(u, C) \in \hat{V}_\exists$ and any edge $(u, v) \in E$ there is exactly one corresponding edge in \tilde{G} : $((u, C), (v, C))$. Given a vertex (u, C) from \tilde{G} , denote by $[(u, C)]_1$ the vertex u . Now, given a strategy $\sigma \in \mathfrak{S}_\exists(\tilde{G})$ we define $s_\sigma \in \mathfrak{S}_\exists(G)$ as follows

$$s_\sigma(v_0v_1v_2\dots) = [\sigma((v_0, C_0)(v_1, C_1 = C_0 \cap E_V(v_0v_1))(v_2, C_1 \cap E_V(v_1v_2))\dots)]_1$$

where $C_0 = E$. It follows from the fact that we have a bijective mapping from plays in \tilde{G} to plays in G which are consistent with positional strategies for Adam, that s_σ is a valid strategy for Eve in G when playing against a positional adversary. Additionally, it is easy to see that s_σ can be realized using finite

memory only. The memory required corresponds to the subsets of E . The current memory element is determined by the applying the operator $E_V(\cdot)$ to the current play prefix.

Now that we have our strategy s_σ for Eve in G , we proceed by proving the analogue of Claim 5 in this setting.

Claim 6. *If $\sigma \in \mathfrak{S}_\exists(\tilde{G})$ is a winning strategy for Eve in \tilde{G} , then*

$$\forall \tau \in \Sigma_V^1(G), \forall i \geq 0 : \mathbf{Val}(\pi_{s_\sigma \tau}[i..]) = v_i \dots \geq \mathbf{cVal}^{v_i}(G \times \tau). \quad (8)$$

Proof. To convince the reader that s_σ has the property from Equation (8), we consider the synchronized product of G and s_σ —that is, the synchronized product of G and the finite Moore machine realizing s_σ . As s_σ is a finite memory strategy, then this product, which we denote in the sequel by $G \times s_\sigma$, is finite. Now, towards a contradiction, suppose that Equation (8) does not hold for s_σ . That is, there is some $\tau \in \Sigma_V^1(G)$ for which the property fails. Further, let us consider an alternative (memoryless) strategy σ' of Eve which ensures $\mathbf{cVal}^v(G \times \tau)$ from all $v \in V$. The latter exists by definition of $\mathbf{cVal}(G \times \tau)$ and memoryless determinacy of discounted-sum games (see, e.g. [15]).

Let H denote a copy of $G \times s_\sigma$ where all edges induced by E from G are added—not just the ones allowed by s_σ —and $H|_{\sigma'}$ denote the sub-graph of H where only edges allowed by σ' are left. Intuitively, both $G \times s_\sigma$ and $H|_{\sigma'}$ are sub-structures of \tilde{G} with a weight function \tilde{w} lifted from w to the blown-up vertex set \tilde{V} . This is due to the way in which we constructed s_σ .

Since, by assumption, s_σ does not have the property of Equation (8) then the edges present in at least one vertex from $H|_{\sigma'}$ and $G \times \sigma$ differ. Note that such a vertex (u, C) is necessarily such that $u \in V_\exists$ —and C is a “memory element” from the machine realizing s_σ corresponding to a subset of E obtained via $E_V(\cdot)$. Furthermore, from our definition of a strategy, we know that there is a single outgoing edge from it in both structures. Let us write (u, v) —instead of $((u, C), (v, D))$ —for the edge in $G \times s_\sigma$ and (u, v') for the edge in $H|_{\sigma'}$. Recall that s_σ is winning for Eve in \tilde{G} . Thus, we have that $(u, v) \notin \tilde{B} = \{((u, C), (v, D)) \in \tilde{E} : u \in V_\exists \text{ and } \exists \tau' \in \Sigma_V^1(G|_C), w(u, v) + \lambda \mathbf{cVal}^v(G \times \tau') < \mathbf{cVal}_{-v}^u(G \times \tau')\}$. It follows that

$$w(u, v) + \lambda \mathbf{cVal}^v(H \times \tau) \geq \mathbf{cVal}^{v'}(H \times \tau).$$

Thus, the strategy σ'' of Eve which takes (u, v) instead of (u, v') and follows σ' otherwise—indeed, this might mean σ'' is no longer memoryless—also achieves at least $\mathbf{cVal}^u(H \times \tau)$ from u onwards. Notice that this process can be repeated for all vertices in which the two structures differ. Further, since both are finite, it will eventually terminate and yield a strategy of Eve which plays exactly as s_σ and for which, since τ was chosen arbitrarily, Equation (8) holds. Contradiction. \square

It follows immediately that $\mathbf{reg}_{\mathfrak{S}_\exists, \Sigma_V^1}^{s_\sigma}(G) = 0$. Indeed, if we suppose that this is not the case, then there exists a strategy $\sigma' \in \mathfrak{S}_\exists(G)$ such that

$$\exists \tau \in \Sigma_V^1(G) : \mathbf{Val}(s_\sigma, \tau) < \mathbf{Val}(\sigma', \tau).$$

The above directly contradicts Claim 6. \square

B.3 Proof of Theorem 3

In this section we present sufficient modifications to our definitions from Section 3 in order for the techniques used therein to be adapted for this case. Particularly, our notion of regret of a play and the safety game used to decide the existence of regret-free strategies need to take into account the fact that witnessing edges taken by Adam affects previously observed local regrets. That is, we formalize the intuition that alternative plays must also be consistent with the behaviour of Adam that we have witnessed in the current play.

We are now ready to define the regret of a play in a game against a positional adversary. Given a play $\pi = v_0 v_1 \dots$, we let

$$\mathbf{reg}(\pi) := \sup\{\lambda^i(\mathbf{cVal}_{-v_{i+1}}^{v_i}(G|_{E_V(\pi)}) - \mathbf{Val}(\pi[i..])) : v_i \in V_\exists\} \cup \{0\}.$$

Consider now a play prefix $\rho = v_0 \dots v_j$. We let the regret of ρ be

$$\max\{\lambda^i(\mathbf{cVal}_{-v_{i+1}}^{v_i}(G|_{E_V(\rho[i..j])}) - \mathbf{Val}(\rho[i..j])) : 0 \leq i < j \text{ and } v_i \in V_\exists\} \cup \{0\}.$$

We will now re-prove Lemma 6 in the current setting.

Lemma 12. *For any strategy σ of Eve,*

$$\mathbf{reg}_{\mathfrak{S}_{\exists}, \Sigma_{\forall}^1}(G) = \sup\{\mathbf{reg}(\pi) : \pi \text{ is consistent with } \sigma \text{ and some } \tau \in \Sigma_{\forall}^1\}.$$

Proof. Consider any $\sigma, \sigma' \in \mathfrak{S}_{\exists}$ and $\tau \in \Sigma_{\forall}^1$ such that $\pi_{\sigma\tau} \neq \pi_{\sigma'\tau}$. Let us write $\pi_{\sigma\tau} = v_0v_1\dots$ and $\pi_{\sigma'\tau} = v'_0v'_1\dots$ and denote by ℓ the length of the longest common prefix of $\pi_{\sigma\tau}$ and $\pi_{\sigma'\tau}$. We claim that

$$\lambda^\ell(\mathbf{cVal}_{\neg v_{\ell+1}}^{v_\ell}(G|E_{\forall}(\pi_{\sigma\tau})) - \mathbf{Val}(\pi_{\sigma\tau}[\ell..])) \geq \lambda^\ell(\mathbf{Val}(\pi_{\sigma'\tau}[\ell..]) - \mathbf{Val}(\pi_{\sigma\tau}[\ell..])). \quad (9)$$

Indeed, if we assume it is not the case, we then get that

$$\mathbf{cVal}^{v'_{\ell+1}}(G|E_{\forall}(\pi_{\sigma\tau})) < \mathbf{Val}(\pi_{\sigma'\tau}[\ell+1..]).$$

However, recall that $G \times \tau$ is a sub-arena of $G|E_{\forall}(\pi_{\sigma\tau})$. Thus, the co-operative value Eve can obtain in the former, say by playing σ' , must be at most that which she can obtain in the latter. Contradiction.

Note that there is another positional strategy τ' for Adam and a second alternative strategy σ'' for Eve which do give us equality for Equation (9). For this purpose, we choose τ' so that $\tau' \in \Sigma_{\forall}^1(G|E_{\forall}(\pi_{\sigma\tau}))$ —so that $\pi_{\sigma\tau}$ is also consistent with τ' , thus $E_{\forall}(\pi_{\sigma\tau}) = E_{\forall}(\pi_{\sigma\tau'})$ (see Lemma 9)—and also such that

$$\mathbf{cVal}^{v'_{\ell+1}}(G \times \tau') = \mathbf{cVal}^{v'_{\ell+1}}(G|E_{\forall}(\pi_{\sigma\tau})).$$

We choose σ'' so that it follows σ for ℓ turns, goes to v' , and then plays co-operatively with τ' from v' . More formally, let σ'' be a strategy for Eve such that $\pi_{\sigma\tau}[\ell] = \pi_{\sigma''\tau}[\ell]$ and therefore, by choice of τ' , such that $\pi_{\sigma\tau'}[\ell] = \pi_{\sigma''\tau'}[\ell]$ and so that

$$\mathbf{Val}(\pi_{\sigma''\tau'}[\ell..]) = \mathbf{cVal}^{v'_{\ell+1}}(G \times \tau').$$

It follows from Equation (9) and the above arguments that for all $\sigma \in \mathfrak{S}_{\exists}$, if there are $\tau \in \Sigma_{\forall}^1$ and $\sigma' \in \mathfrak{S}_{\exists}$ such that $\pi_{\sigma\tau} \neq \pi_{\sigma'\tau}$ then

$$\sup_{\tau, \sigma' : \pi_{\sigma\tau} \neq \pi_{\sigma'\tau}} \lambda^\ell(\mathbf{Val}(\pi_{\sigma'\tau}[\ell..]) - \mathbf{Val}(\pi_{\sigma\tau}[\ell..])) = \lambda^\ell(\mathbf{cVal}_{\neg v_{\ell+1}}^{v_\ell}(G|E_{\forall}(\pi_{\sigma\tau})) - \mathbf{Val}(\pi_{\sigma\tau})). \quad (10)$$

We are now able to prove the result. That is, for any strategy σ for Eve:

$$\begin{aligned} & \sup\{\mathbf{reg}(\pi) : \pi \text{ is consistent with } \sigma \text{ and some } \tau \in \Sigma_{\forall}^1\} \\ &= \sup_{\tau \in \Sigma_{\forall}^1} \mathbf{reg}(\pi_{\sigma\tau} = v_0v_1\dots) && \text{def. of } \pi_{\sigma\tau} \\ &= \sup_{\tau \in \Sigma_{\forall}^1} \max \left\{ 0, \sup_{\substack{i \geq 0 \\ v_i \in V_{\exists}}} \lambda^i \left(\mathbf{cVal}_{\neg v_{i+1}}^{v_i}(G|E_{\forall}(\pi_{\sigma\tau})) - \mathbf{Val}(\pi_{\sigma\tau}[i..]) \right) \right\} && \text{def. of } \mathbf{reg}(\pi_{\sigma\tau}) \\ &= \sup_{\tau \in \Sigma_{\forall}^1} \max \left\{ 0, \sup_{\sigma' : \pi_{\sigma\tau} \neq \pi_{\sigma'\tau}} \lambda^\ell(\mathbf{Val}(\pi_{\sigma'\tau}[\ell..]) - \mathbf{Val}(\pi_{\sigma\tau}[\ell..])) \right\} && \text{by Eq. (10)} \\ &= \sup_{\tau \in \Sigma_{\forall}^1} \max \left\{ 0, \sup_{\sigma' : \pi_{\sigma\tau} \neq \pi_{\sigma'\tau}} (\mathbf{Val}(\sigma', \tau) - \mathbf{Val}(\sigma, \tau)) \right\} && \text{def. of } \mathbf{Val}(\cdot, \ell) \\ &= \sup_{\tau \in \Sigma_{\forall}^1} \sup_{\sigma' \in \mathfrak{S}_{\exists}} (\mathbf{Val}(\sigma', \tau) - \mathbf{Val}(\sigma, \tau)) && 0 \text{ when } \pi_{\sigma\tau} = \pi_{\sigma'\tau} \end{aligned}$$

as required. \square

We will now state and prove a restricted version of Lemma 7. Intuitively, for a play π , we will not be able to consider a deviation with respect to a prefix of π . Rather, we are forced to take the co-operative value with respect to the set $E_{\forall}(\pi)$ —that is, the edges consistent with any positional strategy Adam might be playing—even after the bound on where the best deviation occurs.




Figure 6: Let ρ denote the play prefix $v_0 \dots v_j$. The alternative play from $v_{i'}$ is better than the one from v_i w.r.t ρ . However, for play π' extending ρ , the alternative play from v_i becomes better than the one from $v_{i'}$ if $\lambda^{i'-i} \mathbf{cVal}^{v_{i'}}_{\neg v_{i'+1}}(G \downarrow E_{\forall}(\pi'))$ is smaller than $\mathbf{cVal}^{v_i}_{\neg v_{i+1}}(G \downarrow E_{\forall}(\pi')) - \mathbf{Val}(\rho[i..i'])$.

Lemma 13. Let π be a play in G and suppose $0 < r \leq \mathbf{reg}(\pi)$. Let

$$N(r) := \lfloor (\log r + \log(1 - \lambda) - \log(2W)) / \log \lambda \rfloor + 1.$$

Then $\mathbf{reg}(\pi)$ is equal to

$$\max_{\substack{0 \leq i < N(r) \\ v_i \in V_{\exists}}} \{ \lambda^i (\mathbf{cVal}^{v_i}_{\neg v_{i+1}}(G \downarrow E_{\forall}(\pi)) - \mathbf{Val}(\pi[i..N(r)])) \} - \lambda^{N(r)} \mathbf{Val}(\pi[N(r)..]).$$

Proof. Observe that $N(r)$ is such that $\frac{2W\lambda^{N(r)}}{1-\lambda} < r$. Hence, we have that for all $i \geq N(r)$ such that $v_i \in V_{\exists}$ it holds that $\lambda^i (\mathbf{cVal}^{v_i}_{\neg v_{i+1}}(G) - \mathbf{Val}(\pi[i..])) \leq \frac{2W\lambda^{N(r)}}{1-\lambda} < r$. Clearly, since $\mathbf{cVal}^{v_i}_{\neg v_{i+1}}(H) \leq \mathbf{cVal}^{v_i}_{\neg v_{i+1}}(G)$ holds for any sub-arena H of G , we have that

$$\lambda^i (\mathbf{cVal}^{v_i}_{\neg v_{i+1}}(G \downarrow E_{\forall}(\pi)) - \mathbf{Val}(\pi[i..])) \leq \frac{2W\lambda^{N(r)}}{1-\lambda} < r.$$

It thus follows that

$$\begin{aligned} \mathbf{reg}(\pi) &= \sup \{ \lambda^i (\mathbf{cVal}^{v_i}_{\neg v_{i+1}}(G \downarrow E_{\forall}(\pi)) - \mathbf{Val}(\pi[i..])) : i \geq 0 \text{ and } v_i \in V_{\exists} \} \\ &= \max_{\substack{0 \leq i < N(r) \\ v_i \in V_{\exists}}} \lambda^i \left(\mathbf{cVal}^{v_i}_{\neg v_{i+1}}(G \downarrow E_{\forall}(\pi)) - \mathbf{Val}(\pi[i..N(r)]) \right) - \lambda^{N(r)} \mathbf{Val}(\pi[N(r)..]) \end{aligned}$$

as required. \square

The main difference between the problem at hand and the one we solved in Section 3 is that, when playing against a positional adversary, information revealed to Eve in the present can affect the best alternatives to her current behaviour. Some definitions are in order. Let $\rho = v_0 \dots v_j$ be a play prefix. The *maximal-regret points* of ρ , denoted by $\mathbf{MRP}(\rho)$, is the set

$$\{ 0 \leq i < j : v_i \in V_{\exists} \text{ and } \lambda^i (\mathbf{cVal}^{v_i}_{\neg v_{i+1}}(G \downarrow E_{\forall}(\rho[..j])) - \mathbf{Val}(\rho[i..j])) = \mathbf{reg}(\rho) \};$$

and the *maximal-regret strategies* of ρ , written $\mathbf{MRS}(\rho)$, is equal to

$$\left\{ \tau \in \Sigma_{\forall}^1(G \downarrow E_{\forall}(\rho[..j])) : \bigvee_{i \in \mathbf{MRP}(\rho)} \mathbf{cVal}^{v_i}_{\neg v_{i+1}}(G \downarrow E_{\forall}(\rho[..j])) = \mathbf{cVal}^{v_i}_{\neg v_{i+1}}(G \times \tau) \right\}.$$




Figure 7: A play π' extending ρ in a way such that $\Sigma_V^1(G \downarrow E_V(\pi')) \cap \mathbf{MRS}(\rho) = \emptyset$ cannot have more regret than a play π extending ρ for which $\Sigma_V^1(G \downarrow E_V(\pi)) \cap \mathbf{MRS}(\rho) \neq \emptyset$ —for ρ longer than $\nu(b_G)$.

The above definitions are meant to capture the intuition that, upon witnessing a new choice of Adam, we can reduce the size of the set of possible positional strategies he could be using. Consider a play prefix ρ . The maximal-regret points of ρ correspond to the positions at which best alternatives to ρ occur. The maximal-regret strategies of ρ is the set of positional strategies of Adam, ρ consistent with them, such that at least one of the best alternatives to ρ is consistent with them. Recall from Lemma 9 (ii) that a play prefix ρ is consistent with a positional strategy $\tau \in \Sigma_V^1(G)$ if and only if $\tau \in \Sigma_V^1(G \downarrow E_V(\rho))$. We can, therefore, think of the set of edges $E_V(\rho)$ as representing the set of all positional strategies for Adam in G that ρ is consistent with, *i.e.* $\{\tau \in \Sigma_V^1(G) : \rho \text{ is consistent with } \tau\}$. Let us write $\Sigma_V^1(G, \rho)$ for the set we just described. Let β be the value of one of the best alternatives to ρ . If $\beta' < \beta$ is the value of one of the best alternatives to ρ' , then we know the best alternatives to ρ are not consistent with any strategy from $\Sigma_V^1(G, \rho')$. Then, according to our definition of maximal-regret strategies, this also means that $\mathbf{MRS}(\rho) \cap \Sigma_V^1(G, \rho') = \emptyset$. The converse is also true.

As an example, consider the situation depicted in Figure 6. If, from v_j , the play π' is obtained and we have that $\Sigma_V^1(G \downarrow E_V(\pi')) \cap \mathbf{MRS}(\rho)$ is empty, then the deviation from $v_{i'}$ might no longer be a best alternative. Indeed, there is no positional strategy of Adam which allows the deviation from $v_{i'}$ to obtain the value we assumed (from just looking at the prefix ρ) and which is also consistent with π' . In order to deal with this, we need some more definitions.

Assume that $\mathbf{Reg}_{\Sigma_V^1, \Sigma_V^1}(G) \geq b_G$. For a play prefix $\rho = v_0 \dots v_n$ with $n \geq N(b_G)$, let us define the value δ_ρ (δ for *drop*) as

$$\min_{\substack{0 \leq i \leq j < N(b_G) \\ \tau, \tau' \in \Sigma_V^1(G \downarrow E_V(\rho))}} \left| \lambda^i \left(\mathbf{cVal}_{\neg v_{i+1}}^{v_i}(G \times \tau) - \mathbf{Val}(\rho[i..j]) \right) - \lambda^j \mathbf{cVal}_{\neg v_{j+1}}^{v_j}(G \times \tau') \right|.$$

Intuitively δ_ρ is the minimal drop of the regret achievable by a better alternative (given the information we can extract from ρ).

The smallest possible drop. Let us derive a universal lower bound on δ_ρ for all ρ of length at least $N(b_G)$. In order to do so we will recall “the shape” of the co-operative value of G . Recall the \mathbf{cVal} in a discounted-sum game can be obtained by supposing Eve controls all vertices and computing \mathbf{aVal} instead. It then follows from positional determinacy of discounted-sum games that the \mathbf{cVal} is achieved by a *lasso* in the arena G . More formally, we know that there is a play π in G of the form

$$\pi = v_0 \dots v_{k-1} (v_k \dots v_\ell)^\omega$$

where $0 \leq k < \ell \leq |V|$, and such that $\mathbf{Val}(\pi) = \mathbf{cVal}^{v_0}(G)$. Let us write $\lambda = \frac{\alpha}{\beta}$ with $\alpha, \beta \in \mathbb{Z}$. One can then verify that

Lemma 14. For all sub-arenas H of G , for all vertices $v \in V$, there exists $N \in \mathbb{Z}$ such that $\mathbf{cVal}^v(H) = \frac{N}{D}$ where $D := \beta^{|V|}(\beta^{|V|} - \alpha^{|V|})$.

It then follows from the definition of δ_ρ that:

Lemma 15. For all play prefixes $\rho = v_0 \dots v_n$ such that $n \geq N(b_G)$ we have that

$$\delta_\rho > \frac{1}{\beta^{N(b_G)} D}.$$

Formalizing our claims. We can now prove a replacement for Lemma 7 holds in this context.

Lemma 16. Let π be a play in G and assume $\mathbf{Reg}_{\mathfrak{S}_\exists, \Sigma_\forall^1}(G) > 0$. Let $\nu(b_G)$ denote the value

$$N(b_G) + \left\lfloor \frac{\log(1 - \lambda) - \log W - (N(b_G) + |V|) \log \beta - \log(\beta^{|V|} - \alpha^{|V|})}{\log \lambda} \right\rfloor + 1.$$

Then for all $\sigma \in \mathfrak{S}_\exists$,

$$\sup_{\tau \in \Sigma_\forall^1} \mathbf{reg}(\pi_{\sigma\tau}) = \sup_{\tau \in \Sigma_\forall^1} \mathbf{reg}(\pi_{\sigma\tau}[\dots\nu(b_G)]) - \lambda^{\nu(b_G)} \mathbf{Val}(\pi_{\sigma\tau}[\nu(b_G)\dots]).$$

Proof. Let us consider throughout this argument an arbitrary $\sigma \in \mathfrak{S}_\exists$. From Lemma 13 and the fact that $\nu(b_G)$ is such that $N(b_G)$, we know that $\sup_{\tau \in \Sigma_\forall^1} \mathbf{reg}(\pi_{\sigma\tau} = v_0 \dots)$ equals

$$\sup_{\tau \in \Sigma_\forall^1} \max_{\substack{0 \leq i < \nu(b_G) \\ v_i \in V_\exists}} \{ \lambda^i (\mathbf{cVal}_{\neg v_{i+1}}^{v_i}(G|E_\forall(\pi_{\sigma\tau})) - \mathbf{Val}(\pi_{\sigma\tau}[i..\nu(b_G)])) \} - \lambda^{\nu(b_G)} \mathbf{Val}(\pi_{\sigma\tau}[\nu(b_G)\dots]).$$

Now, also note that $\nu(b_G)$ was chosen so that

$$\frac{W\lambda^{\nu(b_G)}}{1 - \lambda} < \frac{1}{\beta^{N(b_G)+|V|} D}.$$

Hence, for all $\tau' \in \Sigma_\forall^1$ if we write $\pi_{\sigma\tau'} = v'_0 \dots$, then for all $j \geq \nu(b_G)$ such that $v'_j \in V_\exists$ it holds that

$$-\frac{1}{\beta^{N(b_G)+|V|} D} < \lambda^i \mathbf{Val}(\pi_{\sigma\tau'}[i\dots]) < \frac{1}{\beta^{N(b_G)+|V|} D}.$$

It then follows from Lemma 15 and the definition of $\delta_{\pi_{\sigma\tau'}[\dots\nu(b_G)]}$ that, if there exists $\ell \geq \nu(b_G)$ such that for all $0 \leq k \leq \nu(b_G)$ with $v'_k \in V_\exists$

$$\mathbf{cVal}_{\neg v'_{k+1}}^{v'_k}(G|E_\forall(\pi[\dots\ell])) < \mathbf{cVal}_{\neg v'_{k+1}}^{v'_k}(G|E_\forall(\pi[\dots\nu(b_G)]))$$

then $\mathbf{reg}(\pi_{\sigma\tau'}) < \mathbf{reg}(\pi_{\sigma\tau''})$ for all $\tau'' \in \mathbf{MRS}(\pi'[\dots\nu(b_G)])$. This is due to the fact that that $\pi_{\sigma\tau''}[\dots\nu(b_G)] = \pi_{\sigma\tau'}[\dots\nu(b_G)]$ and

$$\mathbf{cVal}_{\neg v'_{k+1}}^{v'_k}(G \times \tau'') = \mathbf{cVal}_{\neg v'_{k+1}}^{v'_k}(G|E_\forall(\pi_{\sigma\tau''}[\dots\nu(b_G)])).$$

The above implies that for all $\sigma \in \mathfrak{S}_\forall$ the value $\sup_{\tau \in \Sigma_\forall^1} \mathbf{reg}(\pi_{\sigma\tau} = v_0 \dots)$ equals

$$\max\{ \mathbf{cVal}_{\neg v_{i+1}}^{v_i}(G|E_\forall(\pi_{\sigma\tau}[\dots\nu(b_G)])) - \lambda^{\nu(b_G)} \mathbf{Val}(\pi_{\sigma\tau}[\dots\nu(b_G)]) : 0 \leq i \leq N(b_G) \text{ and } v_i \in V_\exists \}$$

and therefore (by definition of regret of a prefix) we have that

$$\sup_{\tau \in \Sigma_\forall^1} \mathbf{reg}(\pi_{\sigma\tau}) = \sup_{\tau \in \Sigma_\forall^1} \mathbf{reg}(\pi_{\sigma\tau}[\dots\nu(b_G)]) - \lambda^{\nu(b_G)} \mathbf{Val}(\pi_{\sigma\tau}[\nu(b_G)\dots]).$$

as required. \square

Putting everything together. Let us go back to our example to illustrate how to use $\nu(b_G)$ and the drop of a prefix. Consider now the situation from Figure 7. Recall we have assumed π' is a play extending ρ with $\Sigma_V^1(G|E_V(\pi')) \cap \mathbf{MRS}(\rho) = \emptyset$. It follows that all best alternatives to π' achieve a payoff strictly smaller than $\mathbf{cVal}_{-v_{i'+1}}^{v'}(G|E_V(\rho))$. Thus, the regret of π' can only be bigger than the regret of a play π with $\Sigma_V^1(G|E_V(\pi)) \cap \mathbf{MRS}(\rho) \neq \emptyset$ if the minimal index $k > j$ such that $\Sigma_V^1(G|E_V(\pi'[..j])) \cap \mathbf{MRS}(\rho) = \emptyset$ —*i.e.* the turn at which Adam revealed he was not playing a strategy from $\mathbf{MRS}(\rho)$ —is small enough. In other words, the drop in the value of the best alternative has to be compensated by a similar drop in the value obtained by Eve, and the discount factor makes this impossible after some number of turns.

Proposition 3. *If $\mathbf{Reg}_{\Sigma_\exists, \Sigma_V^1}(G) \geq b_G$ then $\mathbf{Reg}_{\Sigma_\exists, \Sigma_V^1}(G)$ is equal to*

$$\inf_{\sigma \in \Sigma_\exists} \sup \{ \mathbf{reg}(\pi[..\nu(b_G)]) - \lambda^{\nu(b_G)} \mathbf{aVal}^{\hat{u}}(\hat{H}) : \pi = v_0 v_1 \dots \text{ cons. with } \sigma \text{ and some } \tau \in \Sigma_V^1 \}$$

where

- $\hat{u} := (v_{\nu(b_G)}, E_V(\pi[..\nu(b_G)]))$ and
- $\hat{H} := \hat{G} \setminus \{((C, u), (D, v)) : \Sigma_V^1(G|D) \cap \mathbf{MRS}(\pi[..\nu(b_G)]) \neq \emptyset\}$.

Proof. First, note that if $\mathbf{Reg}_{\Sigma_\exists, \Sigma_V^1}(G) > 0$ then there cannot be any regret-free strategies for Eve in G when playing against a positional adversary. It then follows from Corollary 3 that $\mathbf{Reg}_{\Sigma_\exists, \Sigma_V^1}(G) \geq b_G$.

Now using Lemma 16 together with the definition of the regret of a play we get that $\mathbf{Reg}_{\Sigma_\exists, \Sigma_V^1}(G)$ is equal to

$$\inf_{\sigma \in \Sigma_\exists} \sup \{ \mathbf{reg}(\pi[..\nu(b_G)]) - \lambda^{\nu(b_G)} \mathbf{Val}(\pi[\nu(b_G)..]) : \pi \text{ cons. } \sigma \text{ and some } \tau \in \Sigma_V^1 \}.$$

Finally, note that it is in the interest of Eve to maximize the value $\lambda^{\nu(b_G)} \mathbf{Val}(\pi[\nu(b_G)..])$ in order to minimize regret. Conversely, Adam tries to minimize the same value with a strategy from $\mathbf{MRS}(\pi[..\nu(b_G)])$: critically, the strategy is such that the prefix $\pi[..\nu(b_G)]$ is consistent with it. Thus, we can replace it by the antagonistic value from $\pi[\nu(b_G)..]$ discounted accordingly. In this setting we also want to force Adam to play a positional strategy which is consistent with deviations before $N(b_G)$ which achieve the assumed regret of the prefix $\pi[..\nu(b_G)]$. More formally, we have

$$\begin{aligned} & \inf_{\sigma \in \Sigma_\exists} \sup_{\tau \in \Sigma_V^1} \mathbf{reg}(\pi_{\sigma\tau}[..\nu(b_G)]) - \lambda^{\nu(b_G)} \mathbf{Val}(\pi_{\sigma\tau}[\nu(b_G)..]) \\ &= \inf_{\substack{\sigma \in \Sigma_\exists \\ \sigma' \in \Sigma_\exists}} \sup_{\substack{\tau \in \Sigma_V^1 \\ \tau' \in \mathbf{MRS}(\pi_{\sigma\tau}[..\nu(b_G)])}} \mathbf{reg}(\pi_{\sigma\tau}[..\nu(b_G)]) - \lambda^{\nu(b_G)} \mathbf{Val}(\sigma', \tau') \\ &= \inf_{\sigma \in \Sigma_\exists} \sup_{\tau \in \Sigma_V} \mathbf{reg}(\pi_{\sigma\tau}[..\nu(b_G)]) + \inf_{\sigma' \in \Sigma_\exists} \sup_{\tau' \in \mathbf{MRS}(\pi_{\sigma\tau}[..\nu(b_G)])} \left(-\lambda^{\nu(b_G)} \mathbf{Val}(\sigma', \tau') \right). \end{aligned}$$

It should be clear that the RHS term of the sum is equivalent to

$$-\lambda^{\nu(b_G)} \mathbf{aVal}^{\hat{u}}(\hat{H})$$

as required. \square

The above result allows us to claim an EXPSPACE algorithm (when λ is not fixed) to compute the regret of a game. As in Section 3, we simulate the game using an alternating machine which halts in at most a pseudo-polynomial number of steps which depends on $\nu(b_G)$ and, in turn, on b_G . After that, we must compute the antagonistic value of \hat{G} . As a first step, however, we compute the safety game \tilde{G} and determine its winner.

Proposition 4. *Computing the regret value of a game, playing against a positional adversary, can be done in time $\mathcal{O}(\max\{|V|(|E| + 1), \nu(b_G)\})$ with an alternating Turing machine.*

The memory requirements for Eve are as follows:

Corollary 5. *Let $\eta := |\Delta|^d$ where $d = \max\{|V|(|E| + 1), \nu(b_G)\}$. It then holds that $\mathbf{Reg}_{\Sigma_\exists, \Sigma_V^1}(G) = \mathbf{Reg}_{\Sigma_\exists, \Sigma_V^1}(G)$.*




Figure 8: Depiction of the reduction from QBF.




Figure 9: Clause gadget for the QBF reduction for clause $x_i \vee \neg x_j \vee x_k$.

B.4 Lower bounds

In the main body of the paper, namely in Section 4, we have claimed that the regret threshold problem is coNP -hard when λ is fixed. The proof of this claim is provided in Appendix B.4.2. In the next section we shall prove the following result which applies for when λ is not fixed.

Lemma 17. *For a discount factor $\lambda \in (0, 1)$, regret threshold $r \in \mathbb{Q}$, and weighted arena G , determining whether $\text{Reg}_{\exists, \Sigma^1_\forall}(G) \triangleleft r$, for $\triangleleft \in \{<, \leq\}$, is PSPACE-hard .*

B.4.1 Proof of Lemma 17

The QSAT PROBLEM asks whether a given fully quantified boolean formula (QBF) is satisfiable. The problem is known to be PSPACE-complete [10]. It is known the result holds even if the formula is assumed to be in conjunctive normal form with three literals per clause (also known as 3-CNF). Therefore, w.l.o.g., we consider an instance of the QSAT PROBLEM to be given in the following form:

$$\exists x_0 \forall x_1 \exists x_2 \dots \Phi(x_0, x_1, \dots, x_m)$$

where Φ is in 3-CNF. Let n be the number of clauses from Φ .

In the sequel we describe how to construct, in polynomial time, a weighted arena in which Eve ensures regret of at most r if and only if the QBF is true.

We first describe the value-choosing part of the game (see Figure 8). V_{\exists} contains vertices for every existentially quantified variable from the QBF and $V \setminus V_{\exists}$ contains vertices for every universally quantified variable. At each of these vertices, there are two outgoing edges with weight 0 corresponding to a choice of truth value for the variable. For the variable x_i vertex, the **true** edge leads to a vertex from which Eve can choose to move to any of the clause gadgets corresponding to clauses where the literal x_i occurs (see dotted incoming edge in Figure 9) or to advance to x_{i+1} . The **false** edge construction is similar. From the vertices encoding the choice of truth value for x_m Eve can either visit the clause gadgets for it or move to a “final” vertex $\Phi \in V_{\exists}$. This final vertex has a self-loop with weight A .

Our reduction works for values of λ , r , A , B , and C such that the following constraints are met:

- (i) $A < B < C$,
- (ii) $\lambda^2 \left(\frac{C}{1-\lambda} \right) - \lambda^{2nm-2} \left(C + \lambda^2 \frac{B}{1-\lambda} \right) < r$,
- (iii) $\lambda^{2nm-2} \left(C + \lambda^2 \frac{B}{1-\lambda} \right) > \lambda^2 \left(\frac{C \frac{1-\lambda^4}{1-\lambda^8}}{1-\lambda^8} \right)$,
- (iv) $\lambda^2 \left(C + \lambda^2 \frac{B}{1-\lambda} \right) - \lambda^{2nm} \left(\frac{A}{1-\lambda} \right) < r$, and
- (v) $\lambda^{2nm-2} \left(\frac{C}{1-\lambda} \right) - \lambda^{2nm} \left(\frac{A}{1-\lambda} \right) \geq r$.

(See below for a sample concrete assignment.)

Value-choosing strategies. To conclude the proof, we describe the strategy of Eve which ensures the desired property if the QBF is satisfiable and a strategy of Adam which ensures the property is falsified otherwise.

Assume the QBF is true. It follows that there is a strategy of the existential player in the QBF game such that for any strategy of the universal player the QBF will be true after they both choose values for the variables. Eve now follows this strategy while visiting all clause gadgets corresponding to occurrences of chosen literals. At every gadget clause she visits she chooses to enter the gadget. If Adam now decides to take the weight C edge, Eve can go to the center-most vertex and obtain a payoff of at least

$$\lambda^{2nm-2} \left(C + \lambda^2 \frac{B}{1-\lambda} \right),$$

with equality holding if Adam helps her at the very last clause visit of the very last variable gadget. In this case, the claim holds by (i). We therefore focus in the case where Adam chooses to take Eve back to the vertex from which she entered the gadget. She can now go to the next clause gadget and repeat. Thus, when the play reaches vertex Φ , Eve must have visited every clause gadget and Adam has chosen to disallow a weight C edge in every gadget. Now Eve can ensure a payoff value of $\lambda^{2nm} \left(\frac{A}{1-\lambda} \right)$ by going to Φ . As she has witnessed that in every clause gadget there is at least one vertex in which Adam is not helping her, alternative strategies might have ensured a payoff of at most $\lambda^2 \left(C + \lambda^2 \frac{B}{1-\lambda} \right)$, by playing to the center of some clause gadget, or

$$\lambda^2 \left(\frac{C \frac{1-\lambda^4}{1-\lambda^8}}{1-\lambda^8} \right)$$

by playing in and out of some adjacent clause gadgets. By (iii), we know it suffices to show that the former is still not enough to make the regret of Eve at least r . Thus, from (iv), we get that her regret is less than r .

Conversely, if the universal player had a winning strategy (or, in other words, the QBF was not satisfiable) then the strategy of Adam consists in following this strategy in choosing values for the variables and taking Eve out of clause gadgets if she ever enters one. If the play arrives at Φ we have that there is at least one clause gadget that was not visited by the play. We note there is an alternative strategy of Eve which, by choosing a different valuation of some variable, reaches this clause gadget and with the help of Adam achieves value of at least $\lambda^{2nm-2} \left(\frac{C}{1-\lambda} \right)$. Hence, by (v), this strategy of Adam ensures regret of at least r . If Eve avoids reaching Φ then she can ensure a value of at most 0, which means an even greater regret for her. \square




Figure 10: Regret gadget for 2-disjoint-paths reduction.

Example assignment. For completeness, we give one assignment of the positive rationals λ , r , A , B , and C which satisfies the inequalities. It will be obvious the chosen values can be encoded into a polynomial number of bits w.r.t. n and m .

We can assume, w.l.o.g., that $2 \leq 2m \leq n$. Intuitively, we want values such that (i) $A < B < C$ and such that the discount factor λ is close enough to 1 so that going to the center of a clause gadget at the end of the value-choosing rounds, is preferable for Adam compared to doing some strange path between adjacent clauses—this is captured by item (iii). A λ which is close to 1 also gives us item (v) from (i). In order to ensure Eve wins if she does visit the center of a clause gadget, we also would like to have $C - A < r\lambda^{-2}(1 - \lambda)$, which would imply items (ii) and (iv) from the inequality list. It is not hard to see that the following assignment satisfies all the inequalities:

- $\lambda := 1 - \frac{1}{2^{n^3}}$,
- $A := 2$,
- $B := 3$,
- $C := 4$, and
- $r := 3(2^{n^6} - 1)$.

B.4.2 Proof of Theorem 5

The 2-DISJOINT-PATHS PROBLEM on directed graphs is known to be NP-complete [8]. We sketch how to translate a given instance of the 2-DISJOINT-PATHS PROBLEM into a weighted arena in which Eve can ensure regret value of 0 if, and only if, the answer to the 2-DISJOINT-PATHS PROBLEM is negative.

Consider a directed graph G and distinct vertex pairs (s_1, t_1) and (s_2, t_2) . W.l.o.g. we assume that for all $i \in \{1, 2\}$: (i) $s_i \neq t_i$, (ii) t_i is reachable from s_i , and (iii) t_i is a sink (i.e. has no outgoing edges). in G . We now describe the changes we apply to G in order to get the underlying graph structure of the weighted arena and then comment on the weight function. Let all vertices from G be Adam vertices and s_1 be the initial vertex. We replace all edges (v, t_1) incident on t_1 by a copy of the gadget shown in Figure 10. Next, we add self-loops on t_1 and t_2 with weights A and B , respectively. Finally, the weights of all remaining edges are 0. Our reduction works for any value of A and B such that

- (i) $\lambda^{|V|} \frac{A}{1-\lambda} > r$, and
- (ii) $\lambda^{|V|} \frac{B}{1-\lambda} - \lambda \frac{A}{1-\lambda} > r$.

For instance, consider $\alpha := \frac{r+1}{\lambda^{|V|}}$. It is easy to verify that setting $A := (1 - \lambda)\alpha$ and $B := (1 - \lambda)\alpha^2$ satisfies the inequalities. Furthermore, A and B are rational numbers which can be represented using a polynomial number of bits w.r.t. $|V|$ and the size of the representation of both λ and r .

We claim that, in this new weighted arena, Eve can ensure a regret value of 0 if in G the vertex pairs (s_1, t_1) and (s_2, t_2) cannot be joined by vertex-disjoint paths. If, on the contrary, there are vertex-disjoint paths joining the pairs of vertices, then Adam can ensure a regret value strictly greater than r . Indeed, we claim that the strategy that minimizes the regret of Eve is the strategy that, in states where Eve has a choice, tells her to go to t_1 .

First, let us prove that this strategy has regret 0 if, and only if, there are no two paths disjoint paths in the graph between the pairs of states (s_1, t_1) , (s_2, t_2) . Assume there are no disjoint paths, then if

Adam chooses to always avoid t_1 then the regret is 0. If t_1 is reached, then the choice of Eve ensures a value of at least $\lambda^{|V|} \frac{A}{1-\lambda}$. The only alternative strategy of Eve is to have chosen to go to s_2 . As there are no disjoint paths, we know that either the path constructed from s_2 by Adam never reaches t_2 , and then the value of the path is 0 and the regret is 0 for Eve or the path constructed from s_2 reaches t_1 again, and so the regret is also equal to 0 since the discount factor ensures the value of this play is lower than the one realized by the current strategy of Eve. Now assume that there are disjoint paths, if Eve would have chosen to put the game in s_2 (instead of choosing t_1) then Adam has a strategy which allows Eve to reach t_2 and get a payoff of at least $\lambda^{|V|} \frac{B}{1-\lambda}$ while she achieves at most $\lambda \frac{A}{1-\lambda}$. From (i) we have that the regret in this case is greater than r .

To conclude the proof, let us show that any other strategy of Eve has a regret greater than 0. Indeed, if Eve decides to go to s_2 (instead of choosing to go to t_1) then Adam can choose to loop on s_2 and the payoff in this case is 0. The regret of Eve is non-zero in this case since she could have achieved at least $\lambda^{|V|} \frac{A}{1-\lambda}$ by going to t_1 . It follows from (ii) that this ensures a regret value greater than r . \square

C Missing Proofs From Section 5

C.1 Proof of Theorem 8

We reduce the problem to determining the winner of a reachability game on an exponentially larger arena. Although the arena is exponentially larger, all paths are only polynomial in length, so the winner can be determined in alternating polynomial time, or equivalently, polynomial space.

The idea of the construction is as follows. Given a discounted-sum automaton \mathcal{A} , we determinize its transitions via a subset construction, to obtain a deterministic, multi-valued discounted-sum automaton $D_{\mathcal{A}}$. Then we decide if Eve is able to simulate, within the regret bound, the $D_{\mathcal{A}}$ on \mathcal{A} for all *finite* words up to a length (polynomially) dependent on ε . If we simulate the automaton for a sufficient number of steps, then any significant gap between the automata will be unrecoverable regardless of future inputs, and we can give a satisfactory answer for the ε -gap regret problem.

More formally, given a discounted-sum automaton $\mathcal{A} = (Q, q_0, A, \delta, w)$, a regret value r and a precision $\varepsilon > 0$, we construct a reachability game $G_{\mathcal{A}}^{\varepsilon}(r)$ as follows. Let

$$N := \left\lfloor \log_{\lambda} \left(\frac{\varepsilon(1-\lambda)}{4W} \right) \right\rfloor + 1,$$

where W is the maximum absolute value weight occurring in \mathcal{A} , so that $\frac{\lambda^N \cdot W}{1-\lambda} < \frac{\varepsilon}{4}$. Let $P = \{\text{DS}_{\lambda}(\pi) : \pi \in Q^* \text{ is a finite run of } \mathcal{A} \text{ with } |\pi| \leq N\}$ denote the (finite) set of possible discounted payoffs of words of length at most N . Let \mathcal{F} be the set of functions $f : Q \rightarrow \mathbb{R} \cup \{\perp\}$, and for $f \in \mathcal{F}$, let $\text{supp}(f) = \{q \in Q : f(q) \neq \perp\}$. Intuitively, each $f \in \mathcal{F}$ represents a weighted subset of Q ($\text{supp}(f)$ being the corresponding unweighted subset), where $f(q)$ for $q \in \text{supp}(f)$ corresponds to the maximal weight over all (consistent) paths ending in q (scaled by a power of λ). Given $f \in \mathcal{F}$ and $\alpha \in A$ the α -successor of f is the function f_{α} defined as:

$$f_{\alpha}(q') := \begin{cases} \max_{\substack{q \in \text{supp}(f) \\ (q, \alpha, q') \in \delta}} \{\lambda^{-1} \cdot f(q) + w(q, \alpha, q')\} & \text{if this set is not empty} \\ \perp & \text{otherwise.} \end{cases}$$

We define $\mathcal{F}_0 = \{f_0\}$ where $f_0(q_0) = 0$ and $f_0(q) = \perp$ for all $q \neq q_0$; and for all $n \geq 0$, we define $\mathcal{F}_{n+1} := \{f_{\alpha} : f \in \mathcal{F} \text{ and } \alpha \in A\}$. For convenience, let $F = \bigcup_{i=0}^N \mathcal{F}_i$ (considered as a disjoint union).

The game $G_{\mathcal{A}}^{\varepsilon}(r) = (V, V_{\exists}, E, v_0, T)$ is defined as follows:

- $V = (Q \times F \times P) \cup (Q \times F \times P \times A)$;
- $V_{\exists} = (Q \times F \times P \times A)$;
- $((q, f, c), (q, f, c, \alpha)) \in E$ for all $q \in Q$, $f \in F \setminus \mathcal{F}_N$, $c \in P$, and $\alpha \in A$;
- $((q, f, c, \alpha), (q', f', c')) \in E$ for all $q, q' \in Q$, $f \in F \setminus \mathcal{F}_N$, $c \in P$, and $\alpha \in A$ such that $(q, \alpha, q') \in \delta$, $f' = f_{\alpha}$, and $c' = c + \lambda \cdot w(q, \alpha, q')$;

- $v_0 = (q_0, f_0, 0)$; and
- $(q, f, c) \in T$ if, and only if, $f \in \mathcal{F}_N$ and $\max_{s \in \text{supp}(f)} \lambda^{N-1} \cdot f(s) \leq c + r + \frac{\varepsilon}{2}$.

We claim that determining the winner of $G_{\mathcal{A}}^\varepsilon(r)$ yields a correct response for the ε -gap promise problem.

Claim 7. *Let $G_{\mathcal{A}}^\varepsilon(r)$ be defined as above. Then:*

- *If Eve wins $G_{\mathcal{A}}^\varepsilon(r)$ then $\mathbf{Reg}_{\Sigma_\exists, \mathfrak{W}_\forall}(\mathcal{A}) \leq r + \varepsilon$, and*
- *if Adam wins $G_{\mathcal{A}}^\varepsilon(r)$ then $\mathbf{Reg}_{\Sigma_\exists, \mathfrak{W}_\forall}(\mathcal{A}) > r$.*

Proof of Claim 7. It is easy to see that a play of $G_{\mathcal{A}}^\varepsilon(r)$ results in Adam choosing a word $w \in A^*$ of length N , and Eve selecting a run, π , of w on \mathcal{A} by resolving non-determinism at each symbol. Further, if the play terminates at (q, f, c) then $c = \text{DS}_\lambda(\pi)$ and, as f contains the maximal weights of all paths (scaled by a power of λ), $\mathcal{A}(w) = \lambda^{N-1}(\max_{s \in \text{supp}(f)} f(s))$. Since $|w| = N$ we have, for any infinite word $w' \in A^\omega$ and for any run, π' , of \mathcal{A} on w' from q , π' :

$$\begin{aligned} |\mathcal{A}(w \cdot w') - \mathcal{A}(w)| &\leq \frac{\lambda^N \cdot W}{1 - \lambda} < \frac{\varepsilon}{4}, \text{ and} \\ |\text{DS}_\lambda(\pi \cdot \pi') - \text{DS}_\lambda(\pi)| &\leq \frac{\lambda^N \cdot W}{1 - \lambda} < \frac{\varepsilon}{4}. \end{aligned}$$

It follows that:

$$(\mathcal{A}(w) - \text{DS}_\lambda(\pi)) - \frac{\varepsilon}{2} < \mathcal{A}(w \cdot w') - \text{DS}_\lambda(\pi \cdot \pi') < (\mathcal{A}(w) - \text{DS}_\lambda(\pi)) + \frac{\varepsilon}{2}. \quad (11)$$

Now suppose Eve wins $G_{\mathcal{A}}^\varepsilon(r)$. Then, for every word w with $|w| = N$, Eve has a strategy σ that construct a run, π , on \mathcal{A} such that $\mathcal{A}(w) \leq \text{DS}_\lambda(\pi) + r + \frac{\varepsilon}{2}$. We extend this strategy to infinite words by playing arbitrarily after the first N symbols. It follows from Equation 11 that for every infinite word \hat{w} , the resulting run, $\hat{\pi}$,

$$\mathcal{A}(\hat{w}) - \text{DS}_\lambda(\hat{\pi}) < (\mathcal{A}(w) - \text{DS}_\lambda(\pi)) + \frac{\varepsilon}{2} \leq r + \varepsilon.$$

Since $\mathbf{reg}_{\mathcal{A}}^\sigma(\Sigma_\exists, \mathfrak{W}_\forall) = \sup_{\hat{w} \in A^\omega} (\mathcal{A}(\hat{w}) - \text{DS}_\lambda(\hat{\pi}))$, we have $\mathbf{Reg}_{\Sigma_\exists, \mathfrak{W}_\forall}(\mathcal{A}) \leq r + \varepsilon$.

Conversely, suppose Adam wins $G_{\mathcal{A}}^\varepsilon(r)$. Then for any strategy of Eve, Adam can construct a word w , with $|w| = N$ such that the run, π , of \mathcal{A} on w determined by Eve's strategy satisfies $\mathcal{A}(w) > \text{DS}_\lambda(\pi) + r + \frac{\varepsilon}{2}$. Again, from Equation 11 it follows that for any infinite word \hat{w} with w as its prefix and any consistent run $\hat{\pi}'$,

$$\mathcal{A}(\hat{w}) - \text{DS}_\lambda(\hat{\pi}') > (\mathcal{A}(w) - \text{DS}_\lambda(\pi)) - \frac{\varepsilon}{2} > r.$$

As this is valid for any strategy of Eve, we have $\mathbf{Reg}_{\Sigma_\exists, \mathfrak{W}_\forall}(\mathcal{A}) > r$ as required. \square

Now every path in $G_{\mathcal{A}}^\varepsilon(r)$ has length at most N , and as the set of successors of a given state can be computed on-the-fly in polynomial time, the winner can be determined in alternating polynomial time. Hence a solution to the ε -gap promise problem is constructible in polynomial space.

C.2 Proof of Theorem 9

Given an instance of the QSAT PROBLEM – a fully quantified boolean formula (QBF) – we construct, in polynomial time, a weighted arena such that the answer to the regret threshold problem is positive if, and only if, the QBF is true. The main idea behind our reduction is to build an arena with two disconnected sub-graphs joined by an initial gadget in which we force Eve to go into a specific sub-arena. In order for her to ensure the regret is not too high she must now make sure all alternative plays in the other part of the arena do not achieve too high values. In the sub-arena where Eve finds herself, we will simulate the choice of values for the boolean variables from the QBF while in the other sub-arena these choices will affect which alternative paths can achieve high discounted-sum values based on the clauses of the QBF. We describe the reduction for \leq . It will be clear how to extend the result to $<$.




Figure 11: Initial gadget used in reduction from QBF.

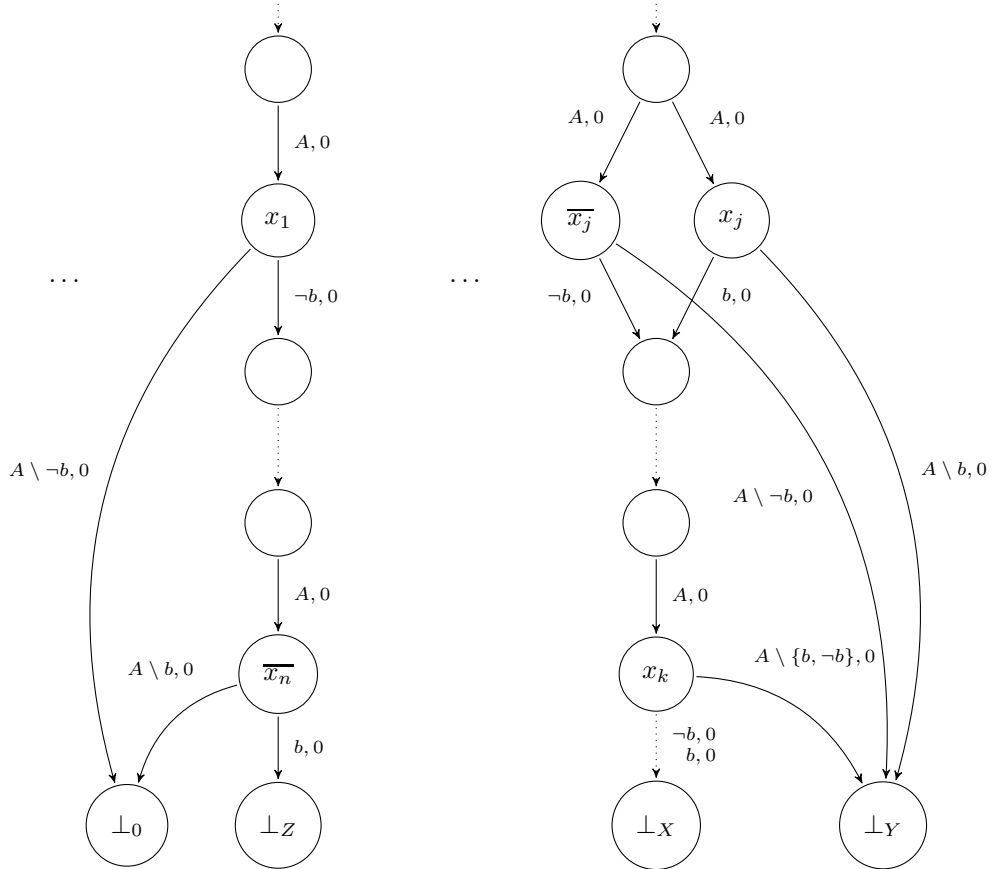


Figure 12: Left and right sub-arenas of the reduction from QBF. Clause i shown on the left; existential and universal gadgets for variables x_j and x_k , respectively, on the right.

The QSAT PROBLEM asks whether a given fully quantified boolean formula (QBF) is satisfiable. The problem is known to be PSPACE-complete [10]. It is known the result holds even if the formula is assumed to be in conjunctive normal form with three literals per clause (also known as 3-CNF). Therefore, w.l.o.g., we consider an instance of the QSAT PROBLEM to be given in the following form:

$$\exists x_0 \forall x_1 \exists x_2 \dots \Phi(x_0, x_1, \dots, x_n)$$

where Φ is in 3-CNF.

We now give the details of the construction. Our reduction works for values of positive rationals r , X , Y , and Z such that

- (i) $\lambda^2 \frac{Z}{1-\lambda} > r + \varepsilon$,
- (ii) $\lambda^{2n} \frac{Z}{1-\lambda} - \lambda^{2n} \frac{X}{1-\lambda} > r + \varepsilon$,
- (iii) $\lambda^{2n} \frac{Z}{1-\lambda} - \lambda^{2n} \frac{Y}{1-\lambda} \leq r$,
- (iv) $\lambda^3 \frac{Y}{1-\lambda} - \lambda^{2n} \frac{X}{1-\lambda} \leq r$.

The alphabet of the new weighted arena is $A = \{bail, b, \neg b\}$.

Example assignment. In order to convince the reader that values which satisfy the above inequalities indeed exist for all possible valuations of n and ε we give such a valuation. Let $f : \mathbb{Q} \rightarrow \mathbb{Q}$ be defined as $f(x) := \frac{(1-\lambda)x}{\lambda^{2n}}$. Note that, w.l.o.g., we can assume that $n \geq 2$. Consider the valuation

- $r := \lambda^{3-2n}(1 + \varepsilon)$,
- $Z := f(r + \varepsilon + 2)$,
- $X := f(1)$,
- $Y := f(2 + \varepsilon)$.

Clearly, inequalities (i)–(iii) hold. Regarding (iv), it will be useful to consider the equivalent inequality

$$\lambda^{3-2n}Y - X \leq \frac{r(1-\lambda)}{\lambda^{2n}}.$$

We observe that the LHS is smaller than $\lambda^{3-2n}(Y - X)$. Furthermore the difference $Y - X$ is equivalent to $\frac{(1+\varepsilon)(1-\lambda)}{\lambda^{2n}}$. Finally, by choice of r we have that the RHS is equivalent to

$$\lambda^{3-2n} \left(\frac{(1+\varepsilon)(1-\lambda)}{\lambda^{2n}} \right).$$

Hence, (iv) holds as well. Note that the chosen values can be encoded into a polynomial number of bits w.r.t. λ and n as well as the size of the representation of ε .

Initial gadget. The weighted arena we construct starts as is shown in Figure 11. Here, Eve has a to make a choice: she can go left or right. If she goes left, then Adam can play *bail* and force her into \perp_0 giving her a value of 0 while an alternative play goes into \perp_Z achieving a value of $\lambda^2 \frac{Z}{1-\lambda}$. By (i) we get that the regret of this strategy is greater than $r + \varepsilon$. Thus, we can assume that Eve will always play to the right.

Choosing values. For each existentially quantified variable x_i we will create a “diamond gadget” to allow Eve to choose a different state depending on the value she wants to assign to x_i . From the corresponding states, Adam will have to play b or $\neg b$, respectively, otherwise he allows her to get to \perp_Y . For universally quantified variables we have a 2-transition path which allows Adam to choose b or $\neg b$ (in the second step). The right path shown in Figure 12 depicts this construction. From (iii) it follows that if Adam cheats at any point during this simulation of value choosing phase of the QSAT game, then the play reaches \perp_Y and the regret is at most r . Hence, we can assume that Adam does not cheat and the play eventually reaches \perp_X . Observe that the choice of values in this gadget is made as follows: at turn $2i$ after having entered the gadget, the value of x_i is decided.

Clause gadgets. For every clause from Φ we create a path in the new weighted arena such that every literal ℓ_i in the clause is synchronized with the turn at which the value of x_i is decided in the value-choosing gadget. That is to say, there are $2i - 1$ states that must be visited before arriving at the state corresponding to ℓ_i . At state ℓ_i , if the value of x_i corresponding to literal ℓ_i is chosen, the play deterministically goes to \perp_0 . Otherwise, traversal of the clause-path continues.

It should be clear that if the QBF is true, then Eve has a value-choosing strategy such that at least one literal from every clause holds. That means that every alternative play in the left sub-arena of our construction has been forced into \perp_0 while Eve has ensured a discounted-sum value of $\lambda^{2n} \frac{X}{1-\lambda}$ by reaching \perp_X . From (iv) it follows that Eve has ensured a regret of at most r . Conversely, if Adam has a value-choosing strategy in the QSAT problem so the QBF is show to be false, then he can use his strategy in the constructed arena so that some alternative path in the left sub-arena eventually reaches \perp_Z . In this case, from (ii) we get that the regret value is greater than $r + \varepsilon$, as expected. \square

C.3 Proof of Theorem 6

Membership. Consider a fixed weighted automaton $\mathcal{A} = (Q, q_I, A, \Delta, w)$ and a discount factor $\lambda \in (0, 1)$. Further, we suppose the regret of \mathcal{A} is 0.

Let us start by defining a set of values which, intuitively, represent lower bounds on the regret Eve can get by resolving the non-determinism of \mathcal{A} on the fly. First, let us introduce some additional notation. Define $\mathcal{A}^q := (Q, q, A, \Delta, w)$, *i.e.* the automaton \mathcal{A} with new initial state q . For states $q, q' \in Q$, let $\mu(q, q') := \sup \left(\{ \mathcal{A}^{q'}(x) - \mathcal{A}^q(x) : x \in A^\omega \} \cup \{0\} \right)$. We are now ready to describe our set of values:

$$M := \{ |w(p, \sigma, q') - w(p, \sigma, q) + \lambda \cdot \mu(q, q')| : p \in Q \text{ and } q, q' \text{ are } \sigma\text{-successors of } p \}.$$

Note that since \mathcal{A} is assumed to be total (*i.e.*, every state-action pair has at least one successor) then M cannot be empty. Observe that, by definition, M only contains non-negative values. Since \mathcal{A} has regret 0, then we know that for all $d \in (0, 1)$, there is a strategy σ_d of Eve such that $\mathbf{reg}_{\mathfrak{S}_{\exists}, \mathfrak{W}_{\forall}}^{\sigma_d}(\mathcal{A}) = 0$. If $M \neq \{0\}$, we let $\varepsilon < \lambda^{|Q|} \cdot (\min M \setminus \{0\})$. Denote by \tilde{Q} the set of states reachable from q_I by reading some finite word x of length at most $|Q|$, *i.e.* $x \in A^{\leq |Q|}$, according to σ_ε . If $M = \{0\}$, let $\tilde{Q} = Q$. We now define a memoryless strategy σ of Eve as follows: if $M = \{0\}$ then σ is arbitrary, otherwise $\sigma(p, a) = q$ implies $q \in \tilde{Q}$. To conclude, we then show that σ ensures regret 0. \square

Hardness. We give a reduction from the SAT problem, *i.e.* satisfiability of a CNF formula. The construction presented is based on a proof in [1]. The idea is simple: given boolean formula Φ in CNF we construct a weighted automaton Γ_Φ such that Eve can ensure regret value of 0 with a positional strategy in Γ_Φ if and only if Φ is satisfiable. Note that this restriction of Eve to positional strategies is no loss of generality. Indeed, we have shown that if the regret of a game against an eloquent adversary is 0, then she has a positional strategy with regret 0.

Let us now fix a boolean formula Φ in CNF with n clauses and m boolean variables x_1, \dots, x_m . The weighted automaton $\Gamma_\Phi = (Q, q_I, A, \Delta, w)$ has alphabet $A = \{bail, \#\} \cup \{i : 1 \leq i \leq n\}$. Γ_Φ includes an initial gadget such as the one depicted in Figure 11. Recall that this gadget forces Eve to play into the right sub-arena. As the left sub-arena of Γ_Φ we attach the gadget depicted in Figure 13. All transitions shown have weight 1 and all missing transitions in order for Γ_Φ to be complete lead to a state \perp_0 with a self-loop on every symbol from A with weight 0. Intuitively, as Eve must go to the right sub-arena then all alternative plays in the left sub-arena correspond to either Adam choosing a clause i and spelling $i\#i$ to reach \perp_1 or reaching \perp_0 by playing any other sequence of symbols. The right sub-arena of the automaton is as shown in Figure 14, where all transitions shown have weight 1 and all missing transitions go to \perp_0 again. Here, from q_0 we have transitions to state x_j with symbol i if the i -th clause contains variable x_j . For every state x_j we have transitions to j_{true} and j_{false} with symbol $\#$. The idea is to allow Eve to choose the truth value of x_j . Finally, every state j_{true} (or j_{false}) has a transition to \perp_1 with symbol i if the literal x_j (resp. $\neg x_j$) appears in the i -th clause.

The argument to show that Eve can ensure regret of 0 if and only if Φ is satisfiable is straightforward. Assume the formula is indeed satisfiable. Assume, also, that Adam chooses $1 \leq i \leq n$ and spells $i\#i$. Since Φ is satisfiable there is a choice of values for x_1, \dots, x_m such that for each clause there must be at least one literal in the i -th clause which makes the clause true. Eve transitions, in the right sub-arena




Figure 13: Clause choosing gadget for the SAT reduction. There are as many paths from top to bottom (\perp_1) as there are clauses (n).




Figure 14: Value choosing gadget for the SAT reduction. Depicted is the configuration for $(x_1 \vee x_2) \wedge (\neg x_1 \vee x_2) \wedge (\neg x_1 \vee \neg x_2)$.

from q_0 to the corresponding value and when Adam plays $\#$ she chooses the correct truth value for the variable. Thus, the play reaches \perp_1 and, as $W = 1$ in the left and right sub-arenas of Γ_Φ , it follows that her regret is 0. Indeed, her payoff will be $\lambda^2/(1 - \lambda)$ —recall the first two turns are spent in the initial gadget, where all transitions leading to both sub-arenas are 0-weighted—which is the maximal payoff obtainable in either sub-arena. If Adam does not play as assumed then we know all plays in Γ_Φ reach \perp_0 and again her regret is 0. Note that this strategy can be realized with a positional strategy by assigning to each x_j the choice of truth value and choosing from q_0 any valid transition for all $1 \leq i \leq n$.

Conversely, if Φ is not satisfiable then for every valuation of variables x_1, \dots, x_m there is at least one clause which is not true. Given any positional strategy of Eve in Γ_Φ we can extract the corresponding valuation of the boolean variables. Now Adam chooses $1 \leq i \leq n$ such that the i -th clause is not satisfied by the assignment. The play will therefore end in \perp_0 while an alternative play in the left sub-arena will reach \perp_1 . Hence the regret of Eve in the game is non-zero. \square